Uncertainty in Climate-Economic Modeling

Svenn Jensen & Christian Traeger*

November 2024

Abstract: This chapter explores the role of uncertainty in climate-economic modeling, particularly in determining the social cost of carbon (SCC) and guiding optimal policy decisions. Various forms of uncertainty–ranging from climate sensitivity and carbon flows to tipping points and economic growth–interact with time and risk preferences, shaping costs, benefits, and urgency of climate policy. Surveying both numerical and analytical approaches, we identify the key drivers of uncertainty's impact on the SCC and climate action. We examine the options and the challenges of integrating these uncertainties into complex climate-economic models, discussing Monte Carlo simulations and models explicitly incorporating decision making under uncertainty. We explore different evaluation frameworks accounting for Epstein-Zin-Weil preferences (disentangled intertemporal risk aversion), ambiguity aversion, and model misspecification. While there is no simple uniform rule how to respond to uncertainty in integrated climate-economic modeling, the field has developed many valuable insights and continues to evolve.

Keywords: Uncertainty, Climate Change, tipping points, ambiguity, Dynamic Programming, Monte Carlo, social cost of carbon, risk premium, Epstein-Zin-Preferences, Integrated assessment, Climate policy

1 Introduction

Uncertainty permeates all aspects of climate change, influencing the physical climate system, the socio-economic system, and our policy response. Uncertainties governing the climate include the critical concept of climate sensitivity—the uncertain response of global temperatures to greenhouse gas emissions (See Figure 3 in Chapter 1). In addition to climate sensitivity, carbon cycle dynamics, including the behavior of carbon sinks and sources, and the possibility of tipping points–such as ice sheet collapse or permafrost thaw–add uncertainty and layers of complexity to the challenge. These physical uncertainties are deeply intertwined with socio-economic uncertainties, such as uncertain economic growth, which significantly affect the optimal level of the social cost of carbon (SCC).

Understanding the nature of uncertainty and its quantification is only the first step in economic models. We must describe not only the uncertainty itself but also its consequences for the decision-

*This paper is written for Volume 1 of the Handbook of the Economics of Climate Change. We are deeply grateful to our colleagues and co-authors, whose insights and collaboration over the years have enriched our understanding of the literature and the role of uncertainty in shaping optimal climate policy. In particular, we thank Larry Karp, David Kelly, Derek Lemoine, Rick van der Ploeg, Christian Gollier, Leo Simon, Marty Weitzman, Antony Millner, Geoff Heal, David Anthoff, Benjamin Crost, Lars Hansen, and Chris Costello. We are also indebted to Lint Barrage and Solomon Hsiang for their valuable feedback on this chapter and their efforts in compiling the handbook. Traeger gratefully acknowledges support from the Research Council of Norway under project 315878. For correspondence, please contact: traeger@uio.no.

making process. Whereas simulations of possible futures can deliver valuable insights, the true challenge of decision making under uncertainty lies in our response. Figure 1 illustrates a decision maker who has to decide how to spread her investments between consumption, physical capital building, and emission reductions under the threat of various climate tipping points. How future decision makers can (or cannot) respond to different events can have crucial implications already for today's optimal decisions and the SCC. We will see that the difference in how we expect future decision makers to respond to realizations of high or low growth or high or low climate sensitivity can, in some cases, even affect the sign of the risk premium. If the reader is willing to take two messages away from this article, then let it be those. First, the policy response to uncertainty is generally not driven by risk aversion. The welfare loss crucially depends on risk aversion, but policies depend on marginal trade-offs, not the level of welfare. Second, it matters whether and how we optimize policies. Simple Monte Carlo simulations are great tools to explore possible futures based on complex models, but they cannot readily analyze how we should respond to uncertainty. And once we analyze optimal policy, it generally matters whether future policy makers respond to the resolution of uncertainty.

Climate change uncertainty differs from a simple objective coin-toss lottery. First, the uncertainty we face is often difficult to translate into probabilities and sometimes merely relies on experts' subjective (gu)estimates. Second, the climate and economic systems are complex, sometimes nonlinear, and feature long delays between actions and impacts. Third, the interaction between multiple uncertainties, such as temperature feedbacks, economic damges, and growth can amplify or mitigate each other in ways that are difficult to predict without formal modeling. Fourth, the way we as a society evaluate these uncertainties through our preferences greatly influences policy recommendations.

The focus of this chapter is on pricing carbon, as it provides the clearest signal across different models on how uncertainty affects the incentives and urgency of climate policy. Rather than concentrating on the level of the SCC itself–which varies widely across models (Moore et al., 2024)–we are primarily interested in how uncertainty alters the SCC. This impact closely mirrors its effect on abatement decisions. Given the limited scope of this chapter, we will not cover how firms and individuals respond to "policy uncertainty," i.e., how governments may set or revise climate policies. While this is an important topic, we anticipate that future volumes will offer dedicated chapters on such second-best settings. For similar reasons, we also abstract from the debate over taxes versus quantities, instead focusing on first-best policy solutions. Even under regulatory delays and uncertainty, price signals or cap-and-trade systems can be adjusted to optimally spread uncertainty between firms and the environment, making the choice between cap-and-trade or carbon taxes less central (Karp and Traeger, 2024).

That said, much of the conceptual analysis in the first sections applies equally well to firms and asset managers, who must consider how future uncertainty should inform their decisions. By exploring how uncertainty shapes the optimal policy response, we provide a framework for understanding the broader implications for both public policy and private decision-making in the face of uncertain climate risks.

A chapter on uncertainty in climate change economics must navigate between providing generic insights on how uncertainty affects decisions and those more specific to climate-economic models. We begin by outlining the general types of uncertainty and discussing approaches for their evaluation further below. Section 2 moves straight into outlining the general structure of current integrated climate-economic models. While more stylized models can offer valuable insights, many of the uncertainties in more structured climate-economic models have more nuanced effects, sometimes altering predictions in ways that overly

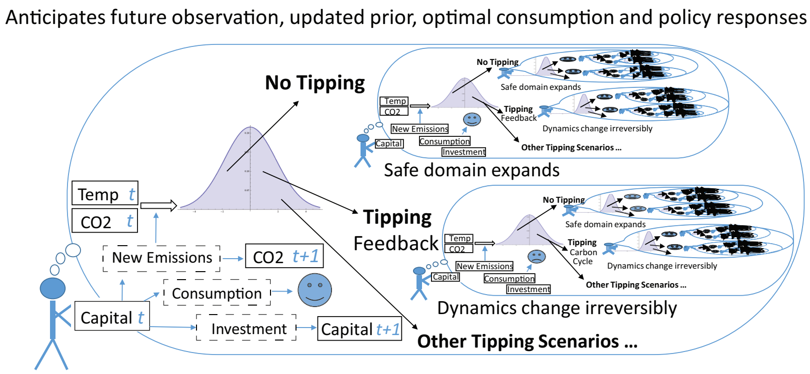

Figure 1: The schematic illustrates how our response to uncertainty today really depends on how we (i) affect future outcomes and (ii) the impact today's choices have while living in different future worlds. The schematic is replicated from Lemoine and Traeger (2016a), who analyze a "tipping point domino". Given tipping either happens or not, the decision tree is still somewhat modest as compared to "simple" growth or climate sensitivity uncertainty, where we we face a continuum of (intertemporally correlated) outcomes in every period. The figure is replicated from Lemoine and Traeger (2016a).

simple models cannot capture. Section 3 discusses general analytic insights, which we consider the most valuable guide for a reader entering this quickly expanding field, where the quantitative estimates for the SCC's risk premium sill vary from negligible to a doubling or tripling of the optimal carbon tax, and also change sign across papers. Section 4 surveys numeric approaches and papers. Again, our emphasis will be less on numbers then on the conceptual approach. The division into analytic and numeric approaches is far from perfect. Our analytic discussions will include comments on magnitudes and signs and some references to numeric approaches, and the papers we discuss in Section 4 will also offer interesting analytic insights and extensions such as the analysis of tipping points and decision making under model uncertainty and robustness.

Interest in the role of uncertainty in climate change has grown rapidly in recent years. While Pindyck (2007) and Heal (2017) general surveys still call for more work on the economic modeling of climate change uncertainty, recent surveys on such uncertainty with a somewhat different focuses include Golub et al. (2014), Rudik and Lemoine (2017), and Cai (2021). We also wish to honor the pioneers of this field, who were ahead of their time. Heal (1984) introduced a stylized analytic assessment of climate change with tipping points 40 years ago. Pizer (1999) applied a log-linear approximation of the DICE model to analyze optimal abatement under uncertainty, finding that uncertainty modestly increases emission reductions. Perhaps most impressively, Kelly and Kolstad (1999) examined a fully stochastic version of the DICE model using neural network approximations of the value function 25 years ago. These early studies have set the stage for contemporary research. We are thrilled to finally see an explosion of the field, including the deployment of many state of the art modeling approaches.

1.1 Fundamental Classifications of Uncertainty

Risk vs Ambiguity (hard uncertainty). One way of classifying different uncertainties is in terms of their probabilistic nature. If a decision maker can assign (unique) probabilities to all events *and* the probabilities are known, then we generally speak of risk. Otherwise, we commonly speak of ambiguity or hard uncertainty. There are many different models of ambiguity. Some assign ranges or sets of probabilities to events, and others replace the concept of a probability distribution with more general concepts. The most common extension of probabilities uses capacities to describe the uncertain world. Capacities are non-additive set functions that are harder to interpret than probabilities. The non-additivity refers to the implication that the probability of two disjoint events no longer has to add up, which implies that the temperature increase by 2100 can be below 2° or it can be above 2° or, well, it can be something else. This difficulty in interpretation explains why this somewhat common model of ambiguity has not found much application in more applied fields such as climate change economics. Yet, certain decision models built on capacities are equivalent to decision models that assign non-unique probabilities to events, a conceptualization of hard uncertainty that is more intuitive and has been prominently applied in the climate change literature.

Objective vs Subjective Probabilities. In situations where we are willing to characterize the world probabilistically, we can interpret probabilities as either objective or subjective. The classical definition of objective probabilities is based on the limit of observing an arbitrarily large frequency of repeated events or on symmetry reasoning (like in the case of a die or a coin). Outside of a casino, such definitions are rarely helpful, and modern probability theory is based on Kolmogoroff's (1933) axioms declaring the existence of a non-negative additive set function normalized to unity where the probabilities of disjoint events add up. These probabilities can be interpreted as (more or less) objective if we have good knowledge of the likelihood of each outcome or as subjective if we do not. Objective probabilities are rare in climate change economics. The best example is the stochasticity of month-to-month or year-to-year temperature fluctuations after detrending, which leads us to the next distinction.

Stochasticity vs Epistemological Uncertainty. A common distinction between uncertainties, which overlaps with the categorizations above, is the following. Stochasticity, or aleatory uncertainty, describes nature's fluctuations of frequent events. We have no way to forecast the exact outcomes because nature's processes are "intrisically stochastic". Given our current forecasting abilities, our best forecast of next month's average temperature will remain the same probabilistic forecast even if we were to observe another ten years of monthly temperature fluctuations. Therefore, stochasticity is also considered an irreducible uncertainty. By contrast, epistemological uncertainty governs a decision maker's lack of knowledge. As we mention in the introduction, scientists are uncertain about the climate's sensitivity to atmospheric CO_2 concentrations. Once we observe the medium-run temperature response to a doubling of atmospheric CO_2, we know the answer and, if still relevant, can predict a similar event deterministically in the future. In this case, nature's process is not stochastic, or at least that's not the most relevant part of the uncerainty. Nature will respond to our added CO_2 in a particular way; we just do not know yet how exactly.

Learning. Connecting immediately to the two classifications above is the idea of learning. Many of the interesting probabilities in climate change economics are neither objective nor do they qualify as stochasticity. As a result, we can learn from observation. The most common modeling approach, which is likely also the most theoretically convincing, is Bayesian learning. It is generally interpreted as combining both subjective uncertainty and objective uncertainty (or stochasticity) in a single framework. In our prime example, we are subjectively uncertain how global average surface temperature responds to a change in atmospheric CO_2 concentrations. Every year, we observe a new temperature response to altered greenhouse gas concentrations. At the same time, temperatures fluctuate naturally from year to year, even without changes in greenhouse gases. This latter stochasticity prevents us from easily learning the temperature response. What we can do in such situations is to follow Bayes rule and update our probabilistic believe over the climate sensitivity based on the observed realization taking into account the known fluctuations of annual temperatures (likelihood function). In order to model Bayesian learning, we have to pick an initial guess for climate sensitivity, the prior. Based on (potentially historic) observations, we use our knowledge of the likelihood function governing annual temperature fluctuations and the particular observation to update the Bayesian prior into a posterior. This posterior will then serve as our new prior. Over time, the prior shrinks. Under some assumptions, it shrinks every period; under other assumptions, it can also increase in response to unexpected observations.

Objective *and* Subjective Probabilities. The idea of combining objective and subjective probabilities links closely to not only Bayesian learning but also a prominent model of (smooth) ambiguity aversion and second-order risk aversion. Here we can associate likelihood functions and objective utilities with objective risk and lower risk aversion and a Bayesian prior with subjective uncertainty that is potentially subject to higher risk aversion as in, e.g., the prominent smooth ambiguity model (Klibanoff et al., 2005, 2009), which we discuss further below and in Section 4.4.

Endogenous vs exogenous uncertainty. It often matters for decisions whether uncertainty is endogenous or exogenous to the decision problem. In our running example, annual temperature fluctuations (deviations from a trend) are mostly exogenous to our anthropogenic doings. By contrast, the uncertainty governing the temperature trend is endogenous to decisions governing emissions or carbon taxation, at least on a global level. Without anthropogenic greenhouse gases, there is no (or far less) uncertainty about the future climate. As we increase our emissions, we increase the uncertainty governing the future.

Ignorance. Ignorance describes a situation where decision makers are unaware of potential outcomes. Even though this could be debated, we include now only the famous “unknown unknowns”, but also the “known unknowns” in this class, where the decision maker is aware that he or she is ignorant about some possible outcomes. Given we do not even conceive these outcomes, they also defy assigning probabilities (or capacities). Arguably, there is a lot of ignorance governing the detailed implications of climate change. Scientists are constantly discovering new implications for ecosystems that defy even the imagination of most of us and clearly that of our models. Yet, by their very nature of not being conceived, it is very hard to build such uncertainties into models. The decision-theoretic literature has built a language to talk about such “known unkowns” and “unknown unknowns” in terms of logical operators (Fagin and Halpern, 1987; Dekel et al., 1998; Heifetz et al., 2013). However, most questions in climate change economics require trade-offs, and trade-offs usually require some form of quantification beyond binary operators or possibilities, which makes it hard to apply these theories in climate change economics.

1.2 Crucial Uncertainties in the Climate and Economic Systems

There are many interacting uncertainties in both the climate and economic systems. One of the most important sources of climate uncertainty is how global temperatures will respond to greenhouse gas emissions. This uncertainty is captured by the so-called *climate sensitivity*, which refers to the long-term increase in global average surface temperature resulting from a doubling of atmospheric CO_2 concentrations. Climate sensitivity is a key determinant of future climate change, as it encompasses a wide range of feedback effects. For example, feedbacks such as water vapor, cloud-formation, ice melting and its impact on our planet's reflectivity, as well as more intricate change such as modifications of the vertical temperature profile in the atmosphere (lapse rate). The IPCC's Sixth Assessment Report estimates the likely range of climate sensitivity to be between 2.5 and 4 degrees Celsius, with a best estimate of 3 degrees Celsius (IPCC, 2021). Some probability distributions still carry substantial wait in the tails making it hard to rule out even substantially stronger temperature responses. Figure 3 in Chapter 1 presents a set of distributions that estimates the climate sensitivity.

Another significant uncertainty in the climate system concerns carbon flows, particularly the future behavior of carbon sinks, such as oceans and terrestrial ecosystems. The ability of these sinks to absorb CO_2 is critical to determining future atmospheric carbon concentrations, which cause the greenhouse effect. The precise mechanics of current carbon flows are still part of active research and the absorptive capacity may decline at higher levels of carbon concentration or temperature, leading to further warming (Denning, 2022).

In recent years, attention has increasingly focused on the potential for tipping points–nonlinear processes that could trigger abrupt and potentially irreversible changes in the climate system. Examples include the collapse of the Greenland ice sheet, the thawing of permafrost, or the dieback of the Amazon rainforest. Each of these events could have cascading impacts on global climate, making tipping points another major source of climate uncertainty (Armstrong McKay et al., 2022). We will focus on the economic response to such threshold-crossing uncertainties in a dedicated Section on tipping points (Section 4.3).

Economic growth and technological advances are another crucial and deeply uncertain component in modeling the interaction between human activity and climate change. They affect both, our emissions that alter the climate system and the impact of climate change feeding back into the economy and human well-being. For example, heatwaves may reduce workers' productivity, sea-level rise may destroy infrastructure, and extreme weather events may cause the loss of agricultural output. Increased average temperatures may make regions uninhabitable and lead to migration or conflict. The extent of these impacts remains highly uncertain, as does the effectiveness of future adaptation efforts. While significant progress has been made in refining the empirical foundations of damage functions, see Chapters 2 and 5, these impacts remain highly uncertain for the foreseeable future (Carleton et al., 2022; Rode et al., 2021).

While some uncertainties are epistemic and may be reduced with better data and models, others are irreducible. The global economy is in constant flux, shaped by innovation, societal changes, and evolving institutions, implying that forecast of economic growth are highly uncertain and will likely remain uncertain (Müller et al., 2022). Policy decisions affect both the climate and the economy and are particularly hard to predict. As climate change is a global externality, mitigating its effects relies heavily on policies enacted at global, regional, and local levels of governance. A key source of uncertainty is how much future decision makers will prioritize climate action by committing to and enforcing effective policies. This uncertainty also extends to how various stakeholders - voters, firms, and governments - will respond to policies or their uncertain anticipation. As we noted above, we will have to save a more careful survey of the response to second-best policy uncertainty to a future volume of the handbook. Section 4 will survey some of the literature that develop future scenario analysis incorporating such reasoning, and we will discuss how the assumed policy response by future policy makers affects today's optimal course of action.

1.3 Uncertainty Evaluation

The question of uncertainty evaluation has several dimensions. Some of the early assessments, as well as some prominent present examples of more complex models, use probabilities for a more structured exploration of possible future scenarios. These scenario explorations prescribe a particular policy and explore the possible consequences. They can also randomly draw policies or policy scenarios making them part of the uncertain future. The climate economic literature generally labels this approach *Monte Carlo analysis* as it is based on relatively straightforward Monte Carlo simulations that are easily implemented in most integrated assessment models (IAMs). We discuss this method in detail in section 4.1. The upside is that we can use any deterministic IAM to sample future worlds. The downside is that there is no uncertainty inside the model itself. Hence, the model cannot properly address how we should respond to uncertainty.

At the other end of the spectrum is a policy maker who optimally responds to uncertainty. She takes decisions keeping in mind a decision tree that quickly branches out over the future. While she might not take future decisions herself, she has to take assumption over how future decision makers will respond to future information, e.g., that growth was lower that expected or the climate sensitivity higher than expected. This forward-looking policy maker also anticipates how today's decisions will influence future outcomes across different possible realizations of the probability distributions. To tackle the quickly unbranching decision tree given the time horizons necessary in a model of climate change the literature resorts to *stochastic dynamic programming* and closely related approaches. The recursive formulation of the optimization makes uncertainty quite tractable and can easily incorporate, e.g., anticipated future learning about the climate sensitivity based on temperature realizations. Yet, the challenge is that numeric dynamic programming struggles with the complexity of many IAMs for reasons we discuss in Section 4.2.

Under either approach, we have to decide how our model evaluates future outcomes, and how much a planner or a representative agent dislikes the uncertainty: we have to specify preferences. This chapter will not delve into the interesting debate of whether climate change should be evaluated normatively or descriptively. We will take the stand that our models should capture observed risk aversion reasonably well. Given we are interested in how decision makers should respond to uncertainty, we emphasize rational attractiveness over behavioral attractiveness. Unfortunately, there are many different perspectives on what constitutes rational behavior. Our main concern is that models do not give rise to foreseeable planning inconsistencies that imply a unnecessary waste of resources (e.g., Dutch books). The vast majority of literature agrees with these criteria, so they do not impose much of a constraint on our discussions.

The standard preference framework for evaluating uncertainty in economics is the *Expected Utility* model. Based on von Neumann and Morgenstern's (1944) aximatization it remains one of the benchmarks for rational decision making. In this model, the curvature of the instantaneous utility function captures the decision maker's attitude toward risk. Most models in climate economics assumes a constant relative risk aversion (CRRA) utility function, which implies decreasing absolute risk aversion as (per capita) consumption increases. Expected utility requires that we quantify uncertainty probabilistically. Thanks to Savage's (1954) alternative axiomatization, we can just as well employ the expected utility framework when probabilities are subjective (e.g., based on expert judgment) or objective (based on known frequencies).

Two criticisms have been raised against the use of expected utility in climate economics. First, in a dynamic setting, the CRRA function simultaneously captures two a priori distinct preference characteristics: the desire for a smooth consumption over time and aversion to risk. Conceptually distinct, they are both captured by the same parameter (or concavity) in the expected utility framework. This constraint can lead to serious difficulties calibrating both risk premia and discount rates to the data, both of which are particularly important in the climate change context. The second critique is that expected utility does not distinguish between different types of uncertainty. Whether a probability distribution reflects objective risk (stochasticity) or expert judgment about epistemic uncertainty, the decision maker treats them the same.

One way to disentangle preferences for consumption smoothing and risk aversion is through *Epstein-Zin-Weil* preferences, which allow for separate functional representations of these two preference dimensions. This formulation enables the decision maker to be intertemporally risk-averse. The use of Epstein-Zin-Weil preferences has become increasingly common in climate economics over the past decade. Epstein-Zin-Weil preferences are recursively defined, requiring a recursive approach to the decision problem, typically solved using stochastic dynamic programming or related recursive methods.

Chapter 1's Figure 3 illustrates how different research groups, relying on various methods, have derived quite different probability distributions for climate sensitivity. From a Bayesian modeling perspective, the IPCC's estimate of the equilibrium climate sensitivity corresponds to a decision maker's subjective prior, and this prior might not even be unique. The subjectivity and multiplicity of distributional estimates raise the question of whether climate sensitivity uncertainty should be treated similarly to known risks or whether we need to treat "deeper uncertainty" differently from objective risk. More specifically, we may need to account for differences in how decision makers respond to ambiguity compared to objective risk, as the two may evoke different levels of aversion.

An extreme form of ambiguity aversion is captured by Gilboa and Schmeidler's (1989)'s *maxmin expected utility*. In this model, the decision maker considers a set of possible probability distributions, evaluates the expected value for each, and then focuses on the worst possible expected outcome. *Robust Control* is a related approach but interprets uncertainty differently, where the decision maker is uncertain whether the model they are using is misspecified (Hansen and Sargent, 2001, 2008).

Klibanoff et al. (2005) introduce a model of *smooth ambiguity aversion*, which they later extend to a recursive dynamic framework in Klibanoff et al. (2009). Smooth ambiguity aversion can account for more moderate attitudes toward ambiguity compared to maxmin expected utility. While the normative implications of ambiguity aversion from a rational decision-making perspective remain a matter of debate, Traeger (2010) shows that smooth ambiguity models are consistent with the von Neumann and Morgenstern (1944) axioms, provided we allow for distinguishing between different types of probabilities or degrees of confidence in probabilistic beliefs.[1]

2 Integrated Climate-Economic Modeling

Section 2.1 summarizes the basic structure of what we consider an integrated climate-economic model. In presenting the model, we do not explicitly highlight uncertainty. In reality, every functional relationship and parameter described is uncertain. For practical purposes, later sections will focus on specific

[1]See Section 4.4 and Berger and Marinacci (2020) for further discussion.

uncertainties. Section 2.2 outlines the basic objective function and explains how uncertainty will be incorporated into the formal model.

Our objective is recursively defined through the Bellman equation, which breaks up long and complex decision trees into a sequence of two-period decision problems. This approach makes uncertainty more tractable but requires the introduction of a value function, which captures the maximum future welfare achievable through optimal management of the climate-economic system. For analytical solutions, the value function provides a useful and insightful tool. For numeric approaches, approximating the value function on a high-dimensional state space can be challenging, as discussed in Section 4.

2.1 General Model Layout

This section briefly lays out the crucial components of what we consider an integrated model of climate and the economy. Chapter 1 discusses a particular example, Nordhaus' DICE model, at some length. Aiming at general insights, we start out with a general model structure. Where needed, we will point out how some contributions rely on more specific assumptions. The present model structure is mainly borrowed from Traeger's (2015,2023) Analytic Climate Economy (ACE) framework. For the moment, we do not pay explicit attention to uncertainty. The reader may think of any of these equations as featuring uncertain functional forms and picking up explicit stochastic parameters. In the later applications, the state variables below will be stochastic, either because the functional form is uncertain or because the inputs themselves are uncertain, e.g., an uncertain temperature increase implies uncertain future production.

Gross output is a function of capital $\boldsymbol{K}_t$, labor $\boldsymbol{N}_t$, energy inputs $\boldsymbol{E}_t$, and the commonly exogenously evolving technology $\boldsymbol{A}_t$.[2] Bold letters indicate that these inputs are generally vectors characterizing the inputs of a multi-sector economy. We assume that climate damages can be characterized as a function of the current atmospheric surface temperature increase above 1900 levels, $T_{1,t}$. Net production is

$$Y_t = F(\boldsymbol{A}_t, \boldsymbol{N}_t, \boldsymbol{K}_t, \boldsymbol{E}_t, T_{1,t}). \tag{1}$$

We assume that the production function $F(\cdot)$ is homogeneous (of arbitrary degree) in capital. This functional form covers the production and abatement structure of several common IAMs such as Golosov et al. (2014) and Nordhaus (2017) as well as a more complex production structure relying on several sectors and energy inputs with distinct and time-changing degrees of energy substitutability Traeger's (2022,2023).

A common assumption is that a multiplicatively separable damage function can capture the impacts of climate change, and the most common functional form is DICE's quadratic damage function where

$$F(\cdot, T_{1,t}) = F^{gross}(\cdot)\frac{1}{1+aT^2} \quad \text{or} \quad F(\cdot, T_{1,t}) = F^{gross}(\cdot)\left(1-aT^2\right).$$

The first version is normalized so that damages do not exceed output.[3] Two calibrations of such damage functions based on Barrage and Nordhaus's (2024) and Howard and Sterner's (2017) are depicted in

[2]Endogenous technological progress with aggregate uncertainty is an interesting challenge to take on in these climate-economic models, which introduces another (set of) equation(s) of motion.

[3]The DICE model versions frequently swap between the different formulations, sometimes even between text and code. Whereas damages could, in principle, exceed damages, they would also quickly after exceed the total capital stock. Of course, we hope we never reach those temperatures.

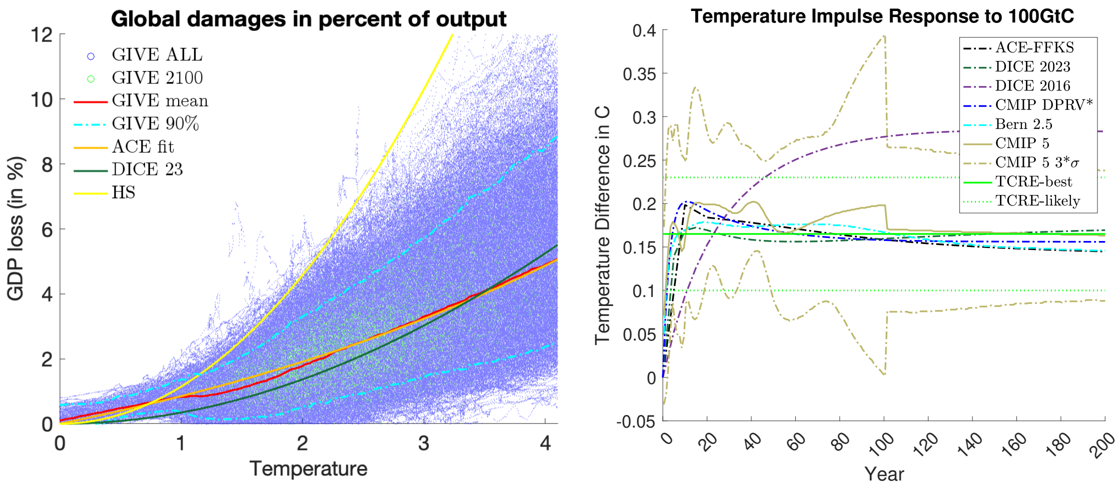

Figure 2: Left panel illustrates the uncertainty governing climate damages based on the GIVE model by Rennert et al. (2022), see Section 4.1. It shows the GDP loss along ten thousand scenarios from the present until 2300 over the temperature increase since 1900. It also shows Traeger's (2021) ACE fit to the mean, using a double-exponential damage function, and the quadratic damage functions of Barrage and Nordhaus's (2024) DICE 2023 model and Howard and Sterner's (2017) preferred estimate from their meta-review. The right panel illustrates the uncertain warming response to today's emissions. It depicts the warming response from emitting a 100Gt carbon pulse into the atmosphere today, based on IPCC's (2013) model comparison study (CMIP 5), showing the mean response and a three-sigma interval around it (Joos et al., 2013). The uncertainty seems to fall 100 years into the future only because some models drop out of the sample for longer horizons. The graph also depicts the impulse response following the approach of Dietz et al. (2021), a ten year time step version of the ACE model (combining a non-linear temperature system with Folini et al.'s (2024) three reservoir carbon cycle), and the DICE 2016 and 2023 models. The TCRE model approximates these responses with a straight line. The Figure is based on Traeger (2021).[5]

Figure 2. The other damage function in the figure is fitted to Rennert et al.'s (2022) GIVE model using the double-exponential (Traeger, 2021)

$$F(\cdot, T_{1,t}) = F^{gross}(\cdot) \exp\Big(-\xi_0 \big(\exp(\xi_1 T_t) - 1\big)\Big). \tag{2}$$

The figure also illustrates the large uncertainty governing the loss of GDP at a given temperature level based on ten thousand simulations by Rennert et al. (2022) and the difference between two prominent quadratic damage estimates (DICE 2023 and HS).

Following Golosov et al. (2014), we assume that the first d entries of the energy vector $\boldsymbol{E}_t$ are fossil-fuel-based and measure their energy content in terms of CO_2. Then emissions are summarized by $E_t \sum_{i=1}^d E_{i,t} + E_t^{other}$, where E_t^{other} captures emissions from land conversion, forestry, and agriculture, often treated as exogenous. Emissions accumulate in the atmosphere where they cause a greenhouse

[5]The right panel also depicts the impulse response of the DICE 2016 model to illustrate some of the recent improvements in connecting climate-economic modeling with the climate sciences. The depicted DICE 2023 impulse response has been modified to align initial conditions to CMIP 5 (Barrage and Nordhaus, 2024). "CMIP DPRV*" follows Dietz et al.'s (2021) approach approximating CMIP 5 using Geoffroy et al.'s (2013) linear approximation of the CMIP 5 results. Based on code kindly provided by the authors, Traeger (2023) adjusted this fit for a 3°C climate sensitivity. ACE's temperature calibration depicted here was not calibration to CMIP 5's impulse experiment but Magicc 6.0 and the IPCC's stochastic transient climate response. We thank Lisa Rennels, Fortunat Joos, David Anthoff, Frank Venmans, and Lint Barrage for their contributions to these graphs.

effect. Over time, some of the added atmospheric CO_2 moves on to other reservoirs, like the oceans or the biosphere (plant matter and soils). Therefore, most models keep track of CO_2 in the different reservoirs, requiring a vector of carbon concentrations that follows the equation of motion

$$\boldsymbol{M}_{t+1} = M(\boldsymbol{M}_t, E_t). \tag{3}$$

In most models, including DICE, equation (3) models a carbon cycle, where the function $M(\dot{})$ is fully characterized by a carbon transition equation $\boldsymbol{\Phi}$, i.e., $\boldsymbol{M}_{t+1} = \boldsymbol{\Phi}\boldsymbol{M}_t + \boldsymbol{e}_1 E_t$.[6] Another popular model, based on Joos et al. (2013), uses several boxes merely for atmospheric carbon. These boxes decay at different rates and mimick the atmospheric impulse response of large climate models.

Increased atmospheric carbon causes radiative forcing, modeled by

$$Q_t = \eta \frac{\log \frac{M_{1,t}+G_t}{M_{pre}}}{\log 2}. \tag{4}$$

where η defines the greenhouse effect's strength and G_t captures the radiative forcing from other, exogenous greenhouse gases in CO_2 equivalents. Radiative forcing is the technical term for the (direct) greenhouse effect, colloquially compared to putting a lid on the pot (which is being heated by the incoming solar radiation not absorbed by CO_2). Fortunately, the oceans keep cooling our planet substantially for the moment. Therefore, most models keep track of at least one more temperature in addition to atmospheric temperature $T_{1,t}$. We denote the equation of motion for the temperature vector by

$$\boldsymbol{T}_{t+1} = T(\boldsymbol{T}_t, Q_t), \tag{5}$$

which closes the model's feedback loop between production and climatic change.

An increasingly popular shortcut for modeling temperature dynamics is based on Matthews et al. (2009), Allen et al. (2009), Dietz et al. (2018) and Dietz and Venmans (2019). The so-called TRCE model (transient climate response to cumulative emissions) simplifies equations (3-5) by assuming that temperatures are linear in cumulative historic emissions. In essence, a ton of CO_2 added to the atmosphere today induces radiative forcing and contributes to atmospheric warming over a long period. Similar to heating a large pot of soup, the warming effect takes time to manifest. Meanwhile, some of the CO_2 moves into the oceans and biosphere. Large climate models show that most of the temperature increase from an additional ton of carbon occurs relatively quickly, within a decade or so, and then remains nearly constant over a long time (as delayed warming balances with slow CO_2 removal). This simplified intuition reflects the "empirical" finding from major climate models that the temperature increase from emissions can, for many purposes, be reasonably approximated by replacing equations (3-5) with this simpler model:

$$\boldsymbol{T}_{1,t+1} = \alpha \sum_{t=\text{preindustrial}}^{t} E_t, \tag{6}$$

and $\alpha \approx 0.165°C$ per 100 Giga tons carbon (GtC) emitted. We currently emit around $10GtC$ per year globally. Thus, every year of emissions warms our planet by approximately $0.017°C$. Figure 2 presents the

[6] The first unit vector $\boldsymbol{e}_1$ implies that we emit into the atmosphere, the first reservoir of the vector $\boldsymbol{M}_t$. Hänsel et al. (2020) extend DICE to include temperature feedbacks into the carbon cycle, an extension now also imported into the latest official DICE update (Barrage and Nordhaus, 2024). In this case, equation (3) generalizes to $\boldsymbol{M}_{t+1} = M(\boldsymbol{M}_t, E_t, T_{1,t})$.

temperature response to a 100GtC emission impulse at today's climate conditions based on Joos et al.'s (2013) model comparison project for the 5^{th} IPCC report (CMIP 5). It depicts the mean response as well as the 3σ interval from these studies (which falls after 100 years because some models left the sample). It also states an approximation to Geoffroy et al.'s (2013) linearization of the model results following Dietz et al.'s (2021) (CMIP DPRV*, see Footnote 5), and a ten year time step illustration based on combining ACE's non-linear temperature system with Folini et al.'s (2024) reference carbon cycle calibration.

On the economic side, we still need to model capital accumulation. Almost all stochastic models assume that there is only one generic capital stock $K_t = \sum_i \boldsymbol{K}_{i,t}$ that is spread across different sectors in every period. It follow the usual equation of motion for capital

$$K_{t+1} = Y_t - C_t + (1 - \delta_k)K_t. \tag{7}$$

Some slightly more complex models such as ACE, or some runs of DICE, also incorporate the potential scarcity of fossil fuels by keeping track of the non-negativity constraint of resource stocks in the ground, which follow the simple equation of motion

$$\boldsymbol{R}_{t+1} = \boldsymbol{R}_t - \boldsymbol{E}_t^d. \tag{8}$$

2.2 Information and Objective Function

For convenience, we now summarize all of our physical state variables in the vector $\boldsymbol{X}_t = (k_t, \boldsymbol{T}_t, \boldsymbol{M}_t, \boldsymbol{R}_t, \boldsymbol{A}_t)$. We note that output Y_t, consumption C_t, and energy use and emissions E_t are flow variables (rather than states). They derive from the state $\boldsymbol{X}_t$ given our policy decisions. When the decision maker does not have full knowledge of the dynamic processes governing the economics and or climate system the dynamic equations are stochastic. Even if only a single equation, say equation (5) for temperature, becomes uncertain, all other equations become stochastic because temperature affects output through equation (1), which then affects all other variables of the model. If shocks are iid, we can merely add a stochastic parameter into any of the equations above to explicitly incorporate the source of the uncertainty. E.g., we can change equation (5) to the form $\boldsymbol{T}_{t+1} = T(\boldsymbol{T}_t, Q_t, \tilde{\epsilon})$, where $\tilde{\epsilon}$ specifies an iid distributed temperature shock.

Most interesting shocks are serially correlated. E.g., annual shocks to global average temperature are highly serially correlated. Similarly, growth shocks are generally serially correlated either because of business cycles or because of long-run growth uncertainty. If shocks are serially correlated, then we need to keep track of past shocks as well, or, we have to add informational state variables to the system that capture past information. In the case of a simple AR(1) shock process, such an informational variable can simply be ϵ_t, i.e., last period's realization of the shock. In a Bayesian learning model governing the climate sensitivity, the informational variables will characterize the distribution of the Bayesian prior. We further discuss such structural learning in Section 4.2. For now, we merely collect such information states by a vector $\boldsymbol{I}_t$, adding another equation of motion[7]

$$\boldsymbol{I}_{t+1} = I(\boldsymbol{I}_t, \boldsymbol{X}_t).$$

[7] Formally, all states and decisions have to be measurable with respect to the filtration generated by the stochastic shocks. Practically, it just means that we cannot base decisions on information we have not received, yet.

Stochastic dynamic programming formulates the objective function recursively. Let us assume that the objective is to maximize the sum of discounted future utility. Let us furthermore assume that the policy maker controls consumption C_t and energy use $\boldsymbol{E}_t$, which by construction includes fossil fuel use and CO_2 emissions, as well as the labor $\boldsymbol{N}_t$ and capital $\boldsymbol{K}_t$ distributions across sectors, which are subject to the obvious availability constraint that we suppress (shares summing to unity). We collect those variables in the control vector $\boldsymbol{a}_t = (C_t, \boldsymbol{E}_t, \boldsymbol{K}_t, \boldsymbol{N}_t)$. Neglecting uncertainty for a moment, we can define the maximal achievable welfare given the current state of the system as

$$V(\boldsymbol{X}_t, \boldsymbol{I}_t) = \max_{\boldsymbol{a}_t} \sum_{t=0}^{\infty} \beta^t u(C_t) \tag{9}$$

where β is the utility discount factor expressing pure time preference, and the optimization is subject to the constraints formulated in equation (1-8). Under certainty, we can cut off the time horizon in a few hundred years and solve this problem solving for consumption and energy inputs of every period, already a daunting task for many sector problems. Under uncertainty, we would have to solve future control variables conditional on future information, which, for most problems, becomes quickly infeasible as the number of control variables explodes. Even if every period only branches out based on two possible stochastic realizations, the number of possible time paths and conditional decision variables in period t evolves with 2^t. If we assume a time step of 10 years, and a binary realization of uncertainty every 10 years, the number of possible time paths 150 years from today is 2^{15}, exceeding 30,000. Counting the number of (information-conditioned) control variables in 150 years, we still have to multiply that number by the number of physical control variables in the model, a bare minimum of 2 in the most aggregated IAMs such as DICE.

We hope this crude reasoning fleshes out the beauty of Bellman's (1966) insight that we can break up the problem above into a sequence of simple two-period problems of the form

$$V(\boldsymbol{X}_t, \boldsymbol{I}_t) = \max_{\boldsymbol{a}_t} \; u(C_t) + \beta \; \mathrm{E}_t \left[V(\boldsymbol{X}_{t+1}, \boldsymbol{I}_{t+1}) \mid \boldsymbol{X}_t, \boldsymbol{I}_t, \boldsymbol{a}_t \right], \tag{10}$$

again subject to equations (1-8). Instead of a quickly unfolding decision tree, the decision maker trades off between current consumption value and future value captured by the value function $V(\boldsymbol{X}_t, \boldsymbol{I}_t)$, which under uncertainty also includes the informational states. The insight is that the problem tomorrow looks exactly the same as the problem today, *conditional* on being in the same state of the world. If the world is changing, we merely add a description of this change explicitly to the state vector.

The crucial difference between solving the Bellman equation (10) and the sum over time (9) is that we only need to solve a single, simple maximization problem, though we do not know the value function $V(\cdot)$ a priori. The key advantage of the Bellman equation is that we no longer have to worry about the branching out of decision trees. Thus, the main challenge becomes finding the value function. If the problem can be solved analytically, the value function immediately provides insight into the shadow values of the various state and information variables. If an analytic solution is not possible, we must approximate the value function, which will be discussed in Section 4.

The Bellman equation (10) explicitly shows that we take the expectation over period $t+1$ using information available in period t. This setting accounts for the fact that future decision makers can incorporate information revealed between today and the time they need to make their decisions. Merely adding an expectation in front of the sum in equation (9) would be somewhat ambiguous as to whether expectations are formed already in period 0 or whether future decisions can incorporate information that is revealed later.

We note that equations (10) and (9) above evaluate the utility of total consumption in the economy. It is more common to maximize per capita consumption, characterizing a representative consumer. Some models also use population-weighted per capita consumption. As noted by Jensen and Traeger (2014), we can usually accommodate labor-weighting and/or per capita consumption by simply renormalizing the discount factor β, which may become time-dependent.[8] For the purpose of this chapter, we abstract from such an induced time dependence of β.

3 Analytic Modeling and Insights

This section derives general analytical insights into the SCC's response to various uncertainties. It integrates key concepts from the literature, drawing in particular from Jensen and Traeger's (2021) "Pricing Climate Risk." We begin with a simple formula for the social cost of carbon (SCC) under uncertainty and proceed to analyze how different uncertainties translate into a risk premium on the optimal carbon tax or SCC. Our primary examples will include uncertainties stemming from economic growth and temperature feedback processes in the climate system, which are generally found to generate the largest risk premiums. Growth uncertainty is particularly interesting for analytical discussion because it can generate both positive and negative risk premiums, depending on the model specification.

The section also briefly relates to Chapter 3 in this handbook and the role of discounting in climate change under uncertainty, discussing some of the implications of using Epstein-Zin-Weil preferences. Most of the section uses general functional forms, which most clearly identify the drivers of risk aversion. As we will see, parameters of specific utility functions can be potentially misleading as they often conflate conceptually different drivers of the risk response. A discussion using general functional forms allows us to understand the economic meaning of the drivers of the uncertainty effect and to sign the risk premium. We close the section with a quantitative model that has a closed-form solution.

3.1 A simple SCC formula

To keep the initial model simple, we summarize the climate system in one equation. The preferred interpretation of this model is the TCRE model following equation (6), where global average surface temperature is proportional to cumulative historic emissions. We keep track of M as cumulative historic emissions rather than T. Alternatively, the model can be interpreted as a simplified model where M denotes atmospheric carbon. Later in Section 3.5, we will revert to this interpretation and explicitly reintroduce the non-linear relation between atmospheric carbon M and temperature dynamics T into the

[8] This is true for using power utility (constant relative risk aversion) or Epstein-Zin-Weil preferences. See, e.g., Traeger (2023) for a discussion of the implications of increasing the weight on more populous future generations. In particular, an often overlooked implication is that, if we interpret the model descriptively, population weighting reduces the rate of pure time preference. If the rate of population growth falls, then population weighting gives a lower rather than higher SCC.

model. In such a model, Jensen and Traeger (2014) show that the SCC is[9]

$$SCC_0 = \mathbf{E}_0 \sum_{t=1}^{\infty} \beta^t \, \frac{u'(c_t)}{u'(c_0)} \underbrace{\left(-\frac{\partial F_t}{\partial M_t}\right)}_{\equiv d(\cdot)} \frac{\partial M_t}{\partial E_0} \; . \tag{11}$$

The SCC captures the impact of a marginal ton of CO_2 released into the atmosphere. Reading the expression from right to left, a marginal increase in emissions affects the cumulative stock of CO_2, which in turn leads to a temperature rise resulting in a marginal loss in production. The production loss causes a proportional change in consumption (see footnote 9), which is then evaluated using marginal utility. The discount factor β^t reflects the utility discount rate applied across time, and we use a lowercase c_t to denote per capital consumption.[10]

Here, uncertainty simply introduces the expected value operator.[11] We abbreviate marginal damages with respect to an increase in CO_2 and, thus, temperature in a given period by $d(\cdot)$, suppressing time indices. The marginal addition of emissions to the stock of CO_2, $\frac{\partial M_t}{\partial E_0}$, is unity in the TCRE model where M_t captures cumulative historic emissions, and this term would fall over time if we were to interpret M_t as the atmospheric CO_2 concentration.

Growth uncertainty plays a particularly interesting (and complex) role for the SCC. Therefore, it helps to briefly review how economic growth affects the SCC under certainty. Let us introduce the common Arrow-Pratt measure of relative risk aversion

$$\text{RRA}(u, c) = -\frac{u''(c)}{u'(c)} c,$$

as the measures of our utility function's curvature. In a deterministic world, it really only expresses an aversion to intertemporal substitution and is the inverse of the elasticity of intertemporal substitution. The constant intertemporal elasticity of substitution (CIES) utility function, $u(c) = \frac{c^{1-\eta}}{1-\eta}$, as used in DICE, gives $\text{RRA}(u, c) = \eta$. Similarly, we define the elasticity of marginal damages with respect to changes in production as

$$\text{Dam}_1(d, y) = \frac{d'(y)}{d(y)} y \; .$$

It measures the percentage increase in marginal damages when production rises by 1%. In the DICE model discussed in Chapter 1, we have $\text{Dam}_1(d, y) = 1$ because damages are proportional to production. This damage sensitivity term relates closely to what Lemoine (2021) later labels the damage scaling effect.

Jensen and Traeger (2014) show that a growth increase raises the SCC if and only if

$$\text{RRA}(u, c) < \text{Dam}_1(d, y) \qquad \Leftrightarrow \qquad \eta < 1 \quad \text{in DICE} \; . \tag{12}$$

[9] Jensen and Traeger (2014) assume a constant consumption rate in the derivation, just like Dietz et al. (2018) discussed further below. Lemoine (2021), also discussed below, uses a more reduced form model without production where $\frac{\partial c_t}{\partial M_t}$ would replace $\frac{\partial F_t}{\partial M_t} = \frac{\partial y_t}{\partial M_t}$.

[10] Under population-weighted per capita consumption, such as in DICE, the discount factor β acquires a time index, as explained in Jensen and Traeger (2014).

[11] In a simple Monte Carlo analysis, it captures the expectations $\mathbf{E}_0$ in period 0. In a more sophisticated recursive dynamic programming setting, it captures the conditional expectations $\mathbf{E}_t$ in period t, which already incorporate information about uncertainty resolution until period t. This distinction only matters if the decision maker can update the policy according to new information.

Technological growth increases the SCC if, and only if, the rise in physical damages (represented by $\text{Dam}_1(d, y)$) outweighs the decrease in the marginal value of damages caused by growing wealth, which is captured by the aversion to intertemporal substitution ($\text{RRA}(u, c)$).

The fact that lower aversion to intertemporal substitution increases the social cost of carbon in the context of economic growth is often framed in terms of the social discount rate. Specifically, reducing the aversion to intertemporal substitution lowers the consumption discount rate, thereby giving more weight to future damages. Additionally, a greater sensitivity of marginal damages to production further supports a stronger optimal abatement effort.

3.2 Uncertainty–Covariance Decomposition and the Climate Beta

The literature has suggested various approaches to analyzing how uncertainty changes the SCC and optimal abatement while maintaining some analytic tractability. We start with a first step suggested by Lemoine (2021), which allows us to connect to Chapter 3's discounting perspective on climate change. For this purpose, we split up the expectation over a single period's contribution to the SCC using the covariance formula. For the moment, we focus on the covariance of marginal utility and marginal damages, i.e., the terms $u'(c)$ and $d(\cdot)$ in equation (11). Marginal utility evaluates the damages caused by an additional ton of CO_2. When marginal utility is high in scenarios where marginal damages are also high, the two effects reinforce each other. This occurs when damages are severe in situations where the representative agent is relatively poor. Conversely, if marginal utility is low when marginal damages are high, the impact of climate damage on welfare is less severe, as those damages occur in scenarios where we are relatively better off to begin with.

Following the asset pricing literature, we can formalize this idea mathematically using the covariance formula for the joint expectation of marginal utility and marginal damages.

$$\mathbf{E}u'(c)d(\cdot) = \mathbf{E}u'(c)\mathbf{E}d(\cdot) + \mathbf{Cov}(u'(c), d(\cdot)), \tag{13}$$

where we suppress time indices and proportionality constants (including $\frac{\partial M_t}{\partial E_0}$). The covariance term states that the valuation of marginal damages is higher if marginal utility and marginal damages are positively correlated. Because marginal utility is not observable, the asset pricing literature frequently approximates marginal utility by consumption, an approximation that is exact under either quadratic utility or normally distributed uncertainty. Assuming a normal distribution of both marginal utility and marginal damages, Stein's Lemma states that $\mathbf{Cov}(u'(c), d(\cdot)) = \mathbf{E}u''(c)\, \mathbf{Cov}(c, d(\cdot))$. Because marginal utility falls ($u''(c) < 0$) in consumption, the valuation in the consumption-based formulation is higher if consumption and marginal damages are negatively correlated.

Lemoine (2021) labels the second term in equation (13) an "insurance" term, and Dietz et al. (2018) focus on it in the context of a "climate beta", a name motivated by the consumption-based capital asset pricing model (CCAPM) presented in Chapter 3. Apart from the risk premium, the crucial determinant of an asset's value in the CCAPM is its correlation with consumption (or consumption growth).[12] Regressing an asset's return on consumption gives rise to the regression coefficient $\beta^{asset} = \frac{\mathbf{Cov}(c,d(\cdot))}{\mathbb{V}\text{ar}(c)}$, referred to as the asset's beta. In the present context, we think of the asset as a small mitigation project that avoids a ton of CO_2 emissions. Its value equals the stream of avoided marginal damages. Section 3.4 will analyze

[12] The CCAPM can be formulated either in terms of consumption itself or consumption growth, both models are mere reformulations of each other.

this covariance for a climate-economic model, after a brief intermission relating our reasoning to some of the insights from the discounting literature.

3.3 A Note on Discounting

Let us, for a moment, assume that marginal damages $d(\cdot)$ are independent of the state of the world so that we can neglect the uncertainty governing $d(\cdot)$. Adding back the time preference discount factor β and time subscripts, equation (13) simplifies to $\beta^t \mathbf{E} u'(c_t) d(\cdot) = d(\cdot) \beta^t \mathbf{E} u'(c_t)$. The term $\beta^t \mathbf{E} u'(c_t)$ captures the marginal value of a deterministic future consumption unit. This expression is the source of Weitzman's (1998) and Weitzman's (2001) famous publications arguing that we should discount future environmental damages at a falling discount rate, an argument that triggered the UK to implement falling discount rates for cost-benefit analysis (Treasury, 2003) and still enjoys some popularity in the field.

By Jensen's inequality, uncertainty increases an expectation if and only if the argument is convex in the stochastic variable. Here, the argument is marginal utility. Marginal utility is convex in consumption if the decision maker is prudent, a common assumption that assumes a positive third derivative of the utility function (more below). Assuming prudence, the expected value falls as uncertainty grows. Thus, as uncertainty grows over time, uncertainty increases our valuation of a future consumption unit. Faced with such an uncertain future, Weitzman offers us the first deterministic consumption unit to ensure safe baseline consumption. Our marginal willingness to pay for this first safe unit grows large with the uncertainty. We will return to this reasoning further below in the climate context. The crucial insight for the moment is that such reasoning relies on the assumption that the future is uncertain, but our environmental payoff is not. Yet, environmental payoffs are generally uncertain as well, in which case Weitzman's reasoning is too simple. Gollier (2004) points out how a seemingly related argument could be construed to turn Weitzman's argument upside down, arguing for an increasing rather than falling discount rate ("Weitzman-Gollier puzzle"). The full equation (13) shows that, in general, the valuation of uncertainty crucially depends on the correlation between the project's payoff and the evolution of the overall economy, a point emphasized by Traeger (2013).

Equation (11) states the SCC in present value terms. It discounts using pure time preference and, otherwise, values marginal damages based on marginal utility. We chose this approach because it makes the interaction between overall consumption, marginal damages, and the impact of emissions on warming explicit. Connecting to the discount rate more generally, we could define the pure rate of time preference ρ and the average consumption growth rate g_t over the time horizon t to transform our factor $\beta^t \, \frac{u'(c_t)}{u'(c_0)} \approx \exp\big(-(\rho + \mathrm{RRA}\, g_t)t\big)$ to a form explicitly containing the Ramsey formula's discount rate $\rho + \mathrm{RRA}\, g_t$. As explained in Chapter 3, this rate picks up a prudence term under uncertainty, an alternative way to understand Weitzman's argument. Yet, once again, it is important to realize that this consumption discount rate is generally correlated with our mitigation project's payoff (and, in particular, marginal damages). Traeger (2014) shows how to take the mitigation project's payoff correlation into account under standard preferences, Epstein-Zin-Weil preferences, ambiguity aversion, and in the case of an unknown correlation (uniform prior over the correlation coefficient). Some climate-economy simulations employ consumption discount rates straight to marginal damages to calculate the SCC. When running stochastic simulations with these essentially deterministic models (as discussed in Section 4 on Monte Carlo simulations), it becomes clear that we cannot simply apply a uniform consumption discount rate across all scenarios: even in a deterministic model the consumption discount rate should adjust according to the Ramsey equation ($\rho + \text{RRA}\, g_t$). Recognizing this limitation, most models adjust the consumption discount rate to at least account for this (essentially deterministic) variation, which recently was labeled 'endogenous discounting' by some authors. We note that this term may be somewhat misleading, as it more commonly refers to methods that incorporate a hazard rate into the discount rate and where the decision maker is fully aware of the uncertainty (Tsur and Zemel, 2009). Given the correlations, this chapter primarily relies on formulations like equation (11), where all components of future value are captured explicitly, rather than separating them into a price and a correlated stochastic discount rate. If we were to split them this way, the time profile of the discount rate, known as the term structure, can be of interest in the context of climate change valuation (Gollier, 2013; Traeger, 2014; Bansal et al., 2019; Newell et al., 2022; Bauer and Rudebusch, 2023).[13]

3.4 Growth Uncertainty–A Negative Risk Premium?

We now proceed to tie the impact of uncertainty on the SCC more closely to our climate-economy model. It is well known that shocks increase an expectation if and only if an upward shock causes a greater positive deviation in the expression than the negative deviation caused by a downward shock of the same magnitude. To analyze this asymmetry in the impact of shocks, it is necessary to examine the curvature of the expression, which indicates how the marginal response differs between positive and negative shocks. We already used this argument above arguing that the expected value increases the expectation if the underlying argument is convex in the stochastic variable.

Section 3.1 explained that the drivers of the SCC's response to a change in growth are changes in marginal utility and marginal damages. In particular, the final term in equation (11) is not affected by economic growth uncertainty. Measuring the curvature of marginal utility, which values marginal damages, we define the relative prudence measure, following Kimball (1990)

$$\text{Prud}(u, c) = -\frac{u'''(c)}{u''(c)} c.$$

In precautionary savings, prudence characterizes the savings response to income uncertainty. For a CIES utility function, $\text{Prud}(u, c) = 1 + \eta = 3$ is always positive, expressing that decision makers' risk aversion (or desire to smooth consumption over time) falls as they grow wealthier.[14] In the same fashion, we introduce a measure of the convexity of *marginal* damage in production

$$\text{Dam}_2(d, y) = \frac{d''(y)}{d'(y)} y.$$

In the DICE model, where damages are proportional to production, $\text{Dam}_2(d, y) = 0$.

We now calculate the approximate risk premium of the SCC. We define this risk premium as the difference between the SCC under uncertainty and the SCC under certainty, where under certainty, consumption and output are fixed to their expected levels denoted by $\bar{c}$ and $\bar{y}$. This risk premium represents the difference in the SCC obtained from upgrading a deterministic IAM by incorporating

[13] It's important to note that a falling term structure (declining interest rate) does not necessarily lead to a higher SCC, as marginal damages could increase even faster as a result of the same uncertainty, and correlation effects must be considered. The term structure is particularly interesting when relating the evaluation directly to capital markets.

[14] We note that for a risk averse decision maker $u''(c) < 0$, hence risk aversion, which is proportional to $-u''(c)$ falls if $u'''(c) > 0$.

uncertainty, while while calibrating both models to the same expected trends. Following Jensen and Traeger (2014), the second-order terms of a Taylor expansion of the expressions in equation (13) delivers period t's contribution to the SCC's risk premium[15]

$$\text{Prud}(u,c)\,\text{RRA}(u,c) + \text{Dam}_1(d,y)\ \text{Dam}_2(d,y) - \underbrace{2\text{Dam}_1(d,y)\,\text{RRA}(u,c)}_{\propto\ \mathbf{Cov}(u'(c),d(y))}\ . \tag{14}$$

Under a small risk approximation, this contribution from period t to the SCC's risk premium is proportional to the variance[16] and, of course, the impact of emissions on warming ($\frac{\partial M_t}{\partial E_0} \approx 1.65 * 10^{-15\circ}C/GtC$ under the TCRE model identifying T with cumulative historic emissions M). Under the assumptions of our IAM, equation (14) pins down the covariance in equation (13) analytically. Under proper normalization, the sought-for covariance is given by the risk-aversion-weighted elasticity of marginal damages to output. Given an approximately constant consumption rate, RRA measures how production shocks reduce marginal utility, and $\text{Dam}_1(d,y)$ measures how they increase marginal damages. Together they characterize the covariance.

The first two terms in equation (14) characterize prudence and the convexity of marginal damage in output. Both derive from the first term on the right side of equation (13), which also gives rise to the deterministic contribution that gets dropped in the characterization of the risk premium. Before interpreting the equation further, we note that even without a small risk approximation, Jensen's inequality tells us that period t's contribution to the SCC is positive if and only if the right side of equation (14) is positive. Following Jensen and Traeger's (2014) rearrangement of the equation, we find that growth uncertainty increases period t's contribution to the SCC if and only if

$$\text{Prud}(u,c) > 2\ \text{Dam}_1(d,y)\ \ - \frac{\text{Dam}_1(d,y)}{\text{RRA}(u,c)}\ \text{Dam}_2(d,y)\ . \tag{15}$$

In a model where damages are proportional to output, $\text{Dam}_2(d,y) = 0$ and we can neglect the final contribution. Then, uncertainty causes a positive risk premium if and only if prudence dominates twice the sensitivity of marginal damages to output. The prudence term results from the expectation $\mathbf{E}u'(c)$ and we discussed in Section 3.3 how it captures an increase in the valuation of marginal damages (and consumption in general) under uncertainty. Yet, we also have to account for the fact that growth uncertainty also translates into the uncertainty of marginal damages, where $\text{Dam}_1(d,y)$ captures that higher output also scales up marginal damages. Under CRRA preferences, $\text{Prud}(u,c) = 1+\eta$. In models like DICE and ACE where damages are proportional to output $\text{Dam}_1(d,y) = 1$ (and $\text{Dam}_2(d,y) = 0$). Then, condition (15) turns into $\eta > 1$, which happens to be the exact opposite of equation (12): growth uncertainty reduces the SCC whenever deterministic growth increases the SCC. Even though these two equations look the same under the assumption of CRRA utility, their economic intuition differs. The uncertainty effect is driven by prudence, and the deterministic effect is driven by a desire for consumption smoothing (measured by RRA).

In models where damages scale non-linearly in production, we also find a contribution to the risk premium from the convexity of marginal damages in output (final term in equations 14 and 15). If the

[15]The first-order terms of the Taylor expansion merely deliver the deterministic contribution, which we subtract again to obtain the risk premium.

[16]Jensen and Traeger (2014) and Dietz et al. (2018) assume proportionality of production and consumption, thus all terms share the same variance that can either be expressed in terms of output or consumption. Lemoine (2021) only models consumption.

damage done by the marginal cost of CO_2 increases more strongly for a positive growth shock than it falls for a negative growth shock, $\text{Dam}_2(d, y)$ further increases the risk premium. In equation (15), this term is scaled by the ratio of the damage elasticity to production and the desire to smooth consumption over time. By equation (12), this ratio is larger than one, amplifying the risk premium if and only if growth itself increases the SCC.

Thus far, our discussion of equation (11) has focused on the contribution to the SCC's risk premium from a single period. Yet, past shocks also affect future periods' marginal damages, particularly if shocks are serially correlated. As a result, serially correlated shocks amplify the risk premium. This effect is particularly prominent in, e.g., Nordhaus's (2008) analysis. Using a Monte Carlo setting, his is either high or low in all periods, and we have a perfect correlation of uncertainty over time and a strong correlation between the realization of the growth rate and temperatures. The study uses DICE to sample 100 Monte Carlo runs featuring eight different uncertainties. Nordhaus finds that growth uncertainty is by far the most important uncertainty in his setting and that uncertainty slightly reduces the SCC. Akin a "climate beta" analysis, Nordhaus shows a positive correlation between high output (low marginal utility) and high temperatures. This now common reasoning is similar to the damage scaling argument embedded in our $\text{Dam}_1(d, y)$, but relies directly on the intertemporal correlation.

What this reasoning overlooks is that higher carbon emissions resulting in higher temperatures do not necessarily increase the marginal damage from another ton of carbon. While the marginal damages increase in temperature, the marginal impact of another ton of CO_2 on warming falls with the prevailing concentration, see equation (4) for radiative forcing. As a result the marginal damages from CO_2 do not necessarily increase in cumulative emissions and temperature (Golosov et al., 2014; Traeger, 2023). Nordhaus does not explicitly analyze the correlation of output with marginal damages, which is done in Dietz et al.'s (2018) more complex Monte Carlo analysis who find a "climate beta" approximately equal to unity.

Dietz et al. (2018) complement their numeric Monte Carlo study with a simple model that addresses intertemporal correlation through a "timeless model," where the same output realization simultaneously drives current consumption and past emissions since the preindustrial era, thereby affecting current temperature. In this parametric example, they demonstrate how incorporating this (admittedly extreme) intertemporal correlation impacts the $\text{Dam}_1(d, y)$ term in equation (14). Focusing on the climate beta, Dietz et al. (2018) do not sign the overall uncertainty effect.

Our derivations above, as well as the Monte Carlo studies, assume that future policy remains fixed. However, it is reasonable to expect that future policymakers will adjust their strategies in response to new information. For example, policymakers can modify emission targets based on whether they are on a high or low growth trajectory. Assuming that future policymakers respond optimally to incoming information, Jensen and Traeger's (2014) solve the forward-looking problem using the Bellman equation (10) and the methods discussed in Section 4.2. They find, in a slightly simplified version of DICE,[17] that under Nordhaus's (2008) parametrization, growth uncertainty increases rather than reduces the SCC, which aligns with the analytical results in equation (15). In addition to incorporating the optimal response of future decision makers, the model also features stochastic growth processes rather than a one-time growth shock. This result is shown in Figure 3, where the dashed line represents $\text{RRA}(u, c) = \eta = 2$, capturing

[17] Jensen and Traeger (2014) use their own version of a simple climate model, rather than DICE's carbon cycle and temperature module. Unlike the TCRE model, which assumes zero decay of cumulative carbon emissions, their model calibrates a small gradually declining decay rate.

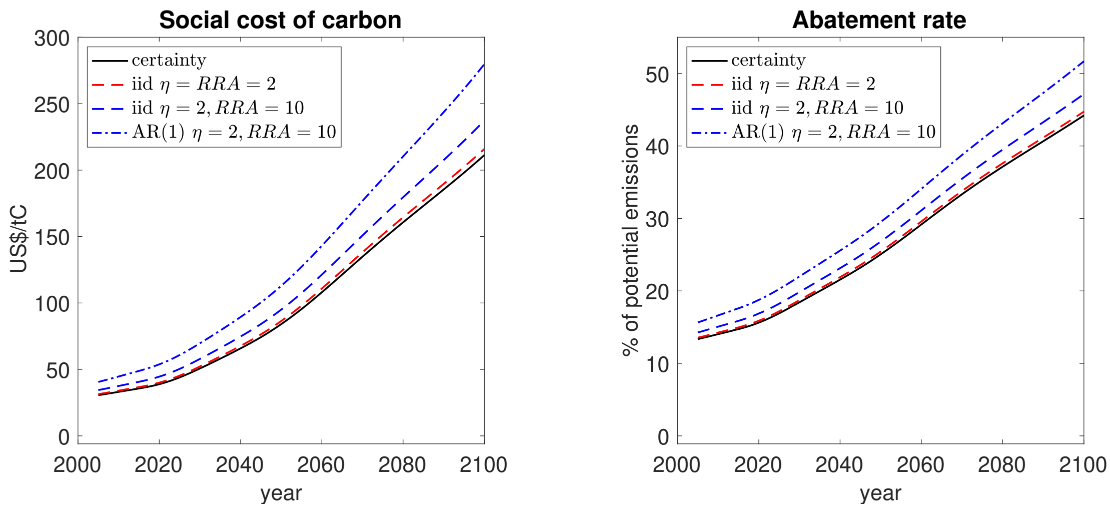

Figure 3: illustrates the effect of uncertainty on the SCC. The red dashed lines represents uncertainty under standard preferences and the blue dashed lines under Epstein-Zin-Weil preferences. The dash-dotted line shows that shock persistence increases the policy's risk premium. All graphs assume an aversion to intertemporal substitution of $\eta = 2$. For the case of Epstein-Zin-Weil the measure of Arrow-Pratt risk averison is RRA $= 10$. For future dates, the paths represent future policy conditional on the expected draws of the growth shocks. The Figure is replicated from Jensen and Traeger (2014).

both the agents' preference for consumption smoothing and their risk aversion. The other graphs increase risk aversion while keeping the preference for consumption smoothing fix, a preference specification we will discuss in Section 3.6, and increase persistence of the growth shock. All graphs depict the future policies conditional on drawing the expected growth realization in every period. These graphs emphasize that even in a world where nature happens to always draw the expected realization, time paths look different based on the uncertainty anticipated by the decision maker.

Finally, we note an alternative decomposition of the expectation in equation (13) by Lemoine (2021). Lemoine uses a parametric setting where damages are proportional to consumption, and utility is isoelastic. In this case, it can make sense to alternatively evaluate the covariance $\mathbf{Cov}\left(cu'(c), \frac{d(\cdot)}{c}\right)$, which shifts all consumption dependence to the left because, by assumption, $\frac{d(\cdot)}{c}$ is independent of consumption. Of course, this rearrangement also affects the overall decomposition in equation (13). Having shifted the consumption dependence out of the marginal damage term, Lemoine finds an additional term that he labels the "damage scaling" term. Using our general functional forms, this damage scaling term precisely picks up our $\text{Dam}_1(d, y)$ contribution above, originally attributed to the "climate beta", leaving the altered covariance term to pick up correlations that are independent of consumption scaling (and absent in our current setting of mere growth uncertainty). In this section, we focused on a simple growth shock that alters output as a mean-zero shock. Appendix A of Jensen and Traeger (2014) relaxes this assumption and shows that alternative growth processes add a term to equations (14) and (15) that accounts for the possible non-linearities in the growth process.

3.5 Climate Uncertainty

The Monte Carlo studies by Nordhaus (2008) and Dietz et al. (2018) both identify temperature uncertainty as the second-most important factor for the SCC's risk premium, after growth uncertainty.

Furthermore, similar to Jensen and Traeger (2013), they find that uncertainty regarding the temperature response to emissions unambiguously increases the SCC.

We follow Jensen and Traeger (2021), explaining how an uncertain climate response changes the reasoning in the previous section somewhat fundamentally. This general analytic reasoning also explains why a simple stochastic process that makes temperatures uncertain, as in the case of growth uncertainty, would overlook crucial contributions to the climate risk premium. More generally, it makes the point that climate change and weather shocks are fundamentally different, not only by nature but also in their contributions to the SCC.

For reasons we further explain in Section 3.7, we abandon the TCRE model in our study of temperature uncertainty and use Jensen and Traeger's (2021) formula for the SCC, which explicitly distinguishes between the impact of emissions on the carbon concentration, $\frac{\partial M_\tau}{\partial E_0}$, and the impact of the atmospheric carbon on atmospheric temperature, $\frac{\partial T_t}{\partial M_\tau}$.[18] Then, the sum in equation (11) has to be replaced by a double sum:

$$SCC_0 = \mathbf{E}_0 \sum_{t=1}^{\infty} \sum_{\tau=1}^{t} \beta^t \frac{u_t'(c_\tau)}{u'(c_0)} \frac{\partial F}{\partial T_t} \frac{\partial T_t}{\partial M_\tau} \frac{\partial M_\tau}{\partial E_0}. \tag{16}$$

An additional ton of carbon today increases atmospheric CO_2 in all future periods giving rise to the sum over τ. In addition, because warming is an autoregressive delay equation, an increase in atmospheric CO_2 and radiative forcing in some period τ causes an increase in temperature in all subsequent periods t.

The response of the average global surface temperature to our emissions is governed by a parameter called the climate sensitivity, see Section 1.2. In a properly calibrated model, it governs both the medium and long-run response to our emissions. A high realization of climate sensitivity increases marginal damages and reduces consumption. These mechanics give rise to a negative correlation of consumption and marginal damages, which is the opposite of the relation discussed above for a growth shock, where a positive growth shock increases consumption while increasing marginal damages (positive correlation).

Jensen and Traeger's (2021) show that this reasoning gives rise to the same formulas as above. Marginal damages are now $d(T) = \frac{\partial F}{\partial T_t}$, where we emphasize temperatures are the crucial argument,[19] and we define

$$\mathrm{Dam}_1(d,T) = \frac{d'(T)}{d(T)} T \quad \text{ and } \quad \mathrm{Dam}_2(d,T) = \frac{d''(T)}{d'(T)} T.$$

In addition, uncertainty now governs climate sensitivity rather than output and consumption directly. As a result, we have to introduce elasticities $\epsilon_{T,s}$ and $\epsilon_{c,s}$ that translate climate sensitivity uncertainty into temperature uncertainty and consumption uncertainty, see Jensen and Traeger (2021) for details. As we reasoned above, the covariance contribution now switches sign, implying that prudence and $\mathrm{Dam}_1(d,T)$ work in the same direction, both increasing the risk premium.[20] In the climate sensitivity context, the

[18] For the purpose of analytic tractability with general function forms, the only endogenous climate states are atmospheric carbon and atmospheric temperature. Ocean cooling is approximated exogenously, and atmospheric carbon is removed at an exogenously fitted time-dependent rate.

[19] Of course, marginal damages depend on all the arguments of the production function. We merely emphasize that we are now analyzing how marginal damages respond to a change in temperature. Here we have $\mathrm{Dam}_1(d,T) = \frac{\partial^2 F_t}{\partial T_t^2} / \frac{\partial F}{\partial T_t}\ T_t$.

[20] The sign of $\mathrm{Dam}_1(d,T)$ stays the same, but the response of consumption to an increase in climate sensitivity is now negative, which can either be captured by defining the elasticities $\epsilon_{T,s}$ and $\epsilon_{c,s}$ with their proper signs (first positive, second

contribution $\mathrm{Dam}_2(d,T)$ captures the convexity of marginal damages, i.e., the change in damage convexity as temperatures increase. Even with damages proportional to output, this term is, in general, no longer zero. While it can initially be positive, it will eventually become negative to ensure that damages cannot outstrip overall production (and/or capital), see Jensen and Traeger (2021) for a discussion.

The major conceptual difference for signing climate uncertainty is, however, the insight that, in the presence of climate change uncertainty, it is no longer enough to analyze the correlation between marginal utility and marginal damages, the reasoning underlying the discussion above and governing equations (13) to (15). Climate sensitivity not only affects the overall temperature and consumption levels, but it also changes how the marginal ton of carbon emitted translates into warming. In equation (16), marginal utility $u'(c_\tau)$, marginal damages $d(T) = -\frac{\partial F}{\partial T_t}$, and marginal warming $\frac{\partial T_t}{\partial M_\tau}$ all depend simultaneously on climate sensitivity. Thus, we need to account for the correlation between all three terms. The climate sensitivity can be defined as the proportionality factor between temperature increase and the radiative forcing from a doubling of CO_2 concentrations. Then, the temperature resulting from a given marginal ton of CO_2 emitted covaries linearly with climate sensitivity. Jensen and Traeger (2021) show that, as a consequence, period t's contribution to the risk premium is proportional to the variance of climate sensitivity and the following sum

$$\underbrace{\mathrm{RRA}\,\epsilon_{c,s}}_{\substack{\text{direct}\\ \text{risk}\\ \text{aversion}}}\big[2+\underbrace{\mathrm{Prud}\,\epsilon_{c,s}}_{\substack{\text{welfare}\\ \text{prudence}}}+\underbrace{3\mathrm{Dam}_1(d,T)\epsilon_{T,s}}_{\substack{\text{welfare}\\ \text{economy}\\ \text{interaction}}}\big]+\underbrace{\mathrm{Dam}_1(d,T)\epsilon_{T,s}}_{\substack{\text{direct}\\ \text{damage}\\ \text{convexity}}}\big[2+\underbrace{\mathrm{Dam}_2(d,T)\epsilon_{T,s}}_{\substack{\text{economy}\\ \text{prudence}}}\big]$$

Strikingly, the risk premium is directly proportional to risk aversion. Usually, risk aversion causes a loss in welfare, but prudence is responsible for the policy response to uncertainty. The reason why risk aversion directly causes a risk premium under climate sensitivity is because of its interaction with climate sensitivity's simultaneous impact on marginal warming (it's a covariance with a linearly changing term). For the same reason, the convexity of the damage function, $\mathrm{Dam}_1(d,T)$, directly increases the risk premium–without interacting with risk aversion as in equation (14). Jensen and Traeger (2021) show that this direct impact of the damage convexity implies by far the biggest risk premium under any model specification. We emphasize that any model featuring climate change uncertainty as an exogenous stochastic process governing temperatures will overlook this most substantial contribution to the SCC under climate uncertainty.

The Monte Carlo assessments by Nordhaus (2008), Dietz et al. (2018), and Lemoine (2021) all pick up the combined effects above and find that temperature uncertainty has a substantial impact, but that the magnitude of the SCC's risk premium is dominated by growth uncertainty.[21] Jensen and Traeger (2021) split up the different uncertainty channels above. They find that the damage convexity channel by far dominates all the other channels. On second place, with approximately equal contributions, are the direct risk aversion effect and the "welfare-economy interaction", a term that captures the interplay of risk aversion and damage convexity and reflects the covariance between marginal utility and marginal damages. For more convex damage functions, such as a cubic-in-temperature specification, "economy prudence" comes in second place, outpacing risk aversion and the welfare-economy interaction. Economy

negative), or by defining the elasticities to be positive and literally switching the sign in front of the $\mathrm{Dam}_1(d,T)$ term in equations (14) and (15).

[21] Using somewhat more extreme uncertainty specifications, Lemoine finds that damage uncertainty has an even stronger impact.

prudence captures the change in damage convexity, just like conventional welfare prudence captures the change in risk aversion. Its contribution reflects that an increasing convexity of the damage function adds to the incentive to cut back on emissions under uncertainty, where convexity reduces expected output in the economy.

Once more, the above reasoning and Monte Carlo studies relied on keeping policies fixed. Jensen and Traeger (2021) show that this assumption exaggerates the risk premium by comparing the formula's output (and the Monte Carlo result) to the solution to the forward-looking problem where policy adapts to (anticipated) learning over climate sensitivity. They also show that the main effects of policy response can be captured by adjusting the (endogenous) temperature elasticity to changes in climate sensitivity $\epsilon_{T,s}$, which has to pick up that optimal policy fights the warming more intensely if climate sensitivity turns out to be high rather than low in the future.

3.6 Epstein-Zin-Weil preferences

Climate change assessments need to account for both time and risk preferences. The well-known risk-premium puzzle in asset pricing demonstrates that standard utility functions struggle to explain observed behavior under risk without significantly increasing risk aversion, which, in turn, leads to inconsistencies with observed risk-free rates. In typical models, raising risk aversion to capture risk premia also intensifies intertemporal consumption smoothing, effectively increasing the discounting of the future. To address this issue, the long-run risk literature emphasizes the importance of distinguishing between risk attitude and intertemporal substitution, as these two factors influence decision-making in different ways.

Empirical evidence supports the idea that individuals are much more averse to risk than they are to intertemporal substitution. The literature estimates the coefficient of relative risk aversion for a representative household to be between 6 and 8 (Vissing-Jørgensen and Attanasio, 2003; Bansal and Yaron, 2004; Bansal et al., 2010, 2012; Chen et al., 2013; Nakamura et al., 2013; Bansal et al., 2014; Collin-Dufresne et al., 2016; Nakamura et al., 2017). Simply increasing the curvature of the utility function to account for such a risk attitude would lead to excessively high risk-free discount rates, a result at odds with market observations (known as the risk-free rate puzzle), and would discount away any concern for climate change in economic models. The assumption that attitude to risk coincides with the aversion to intertemporal consumption change is embedded in the intertemporally additive expected utility model, yet it is not a necessary result of von Neumann and Morgenstern's (1944) axioms (Traeger, 2019). Hence, this section analyzes how the SCC's risk premium changes if we distinguish risk attitude from the propensity to smooth consumption over time. The most prominent example of such preferences are those by Epstein & Zin (1989,1991) and Weil (1990).

Ha-Duong and Treich (2004) analyze Epstein-Zin-Weil preferences in an early two-period pollution control problem, and Crost and Traeger (2014) introduce them into a recursive stochastic DICE-based IAM with infinite planning horizon and stochastic damages. Here, we follow Jensen and Traeger (2014), who incorporate Epstein-Zin-Weil preferences into the growth uncertainty setting we already discussed above. The authors follow Traeger's (2007) reformulation of Epstein-Zin-Weil preferences, rewriting the Bellman equation as

$$V(\boldsymbol{X}_t, \boldsymbol{I}_t) = \max_{\boldsymbol{a}_t} \; u(C_t) + \beta f^{-1} \left(\mathbf{E}_t \left[f V(\boldsymbol{X}_{t+1}, \boldsymbol{I}_{t+1}) \mid \boldsymbol{X}_t, \boldsymbol{I}_t \right] \right) \tag{17}$$

where the concavity of u continues to capture aversion to intertemporal substitution, and the function

f picks up risk aversion above and beyond the mere desire to smooth consumption over time. Traeger (2007,2019) gives an axiomatic characterization of this additional risk aversion as intrinsic or intertemporal risk aversion. In the Epstein-Zin-Weil case, f is a power function. The difference between Traeger's formulation and that of Epstein & Zin (1989,1991) and Weil (1990) is that our Bellman equation above is additive in the time step, which makes it more obvious that u captures consumption smoothing and the generalized mean $f^{-1}(\mathbf{E}_t f[\cdot])$ captures additional aversion to risk (in the deterministic case, where $\mathbf{E}_t$ is degenerate, f and f^{-1} cancel).

Borrowing from Traeger's (2011) formulation in the discounting context, Jensen and Traeger (2014) show that Epstein-Zin-Weil preferences change equation (11) for the SCC under uncertainty to the form

$$SCC_0 \propto \mathbf{E}_0 \sum_{t=0}^{\infty} \beta^t \underbrace{\left\{ \prod_{j=0}^{t-1} \Pi_j P_j \right\}}_{\text{EZW-preferences}} \frac{u'(c_t)}{u'(c_0)} \left(-\frac{\partial F_t}{\partial M_t} \right) \frac{\partial M_t}{\partial E_0} \,. \tag{18}$$

This equation still represents the social cost of carbon as the discounted sum of future marginal damages. The term $\prod_{j=0}^{t-1} \Pi_j P_j$ introduces the additional adjustment for (intertemporal) risk aversion inherent in Epstein-Zin-Weil preferences. In this context, Π_j captures intertemporal prudence, and P_j represents pessimism, which biases the decision maker toward giving more weight to bad outcomes. Under standard preferences with certainty, both Π_j and P_j would equal one. We now explain these contributions in more detail and how they affect the SCC's risk premium.

The intertemporal prudence term Π_t is defined as

$$\Pi_t = \frac{\mathbf{E}_t f'(V_{t+1})}{f'(f^{-1}\mathbf{E}_t f(V_{t+1}))} \,.$$

This term compares the expected marginal value of future welfare, $\mathbf{E}_t f'(V_{t+1})$, to the marginal value at the expected future welfare, $f'(f^{-1}\mathbf{E}_t f(V_{t+1}))$ (where expected future welfare is taken using the generalized mean).

For mean-zero shocks to future welfare, the prudence term usually increases the right-hand side of equation (18), raising the SCC (as also more generally investment under uncertainty). More precisely, it does so if, and only if, absolute intertemporal risk aversion, given by $-\frac{f''}{f'}$, decreases as welfare increases (Traeger, 2011). In simpler terms, this condition means that as future welfare improves, the decision maker becomes less averse to risk, which encourages more precautionary saving. We can express this condition as $\text{Prud}(f, V) > \text{RRA}(f, V)$, meaning that the third derivative of f (representing intertemporal prudence) dominates the second derivative (representing intertemporal risk aversion). In this case, prudence is strong enough to encourage additional saving in response to uncertainty. This condition, known as decreasing absolute intertemporal risk aversion, is satisfied for the standard parametrization of Epstein-Zin-Weil preferences. However, with growth uncertainty, mean-zero technology shocks do not necessarily result in mean-zero welfare shocks. As a result, the overall response of the intertemporal prudence term is slightly more complicated and we refer the interested reader to Jensen and Traeger (2014) for more details.

The quantitatively dominating correction of the SCC under Epstein-Zin-Weil preferences results from pessimism weighting term defined as

$$P_t = \frac{f'(V_{t+1})}{\mathbf{E}_t f'(V_{t+1})} \,.$$

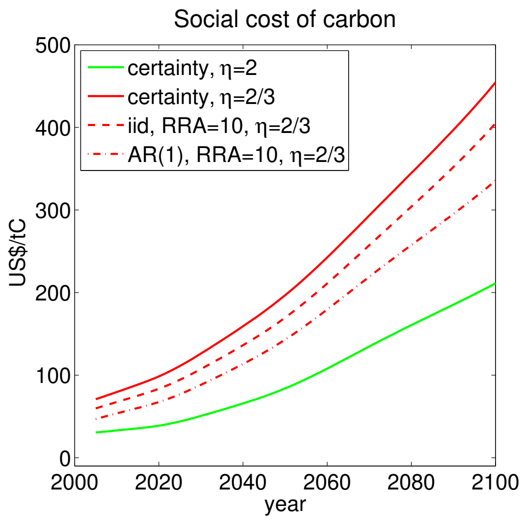

Figure 4: Left panel shows the SCC's response to iid (dashed) and more persistent (dash-dotted) growth shocks in red. Here, aversion to intertemporal substitution is $\eta = 2/3$ and uncertainty causes a negative risk premium. Both panels employ Epstein-Zin-Weil preferences with an Arrow-Pratt coefficient of relative risk aversion of 10. The upward shift from the green to the red solid line under certainty represents the recalibration of the intertemporal consumption smoothing parameter under Epstein-Zin-Weil preferences, which is also reflected in the falling green solid line in the right panel. The right panel shows that the SCC's risk premium is negative under a low aversion to intertemporal substitution η and switches sign for η slightly above unity. The Figure is replicated from Jensen and Traeger (2014).

The term P_t represents a normalized weight referred to as a pessimism bias. This name reflects that, for a concave risk aversion function f, low welfare outcomes receive a higher weight through P_t, while high welfare outcomes receive a lower weight. In other words, the decision maker effectively places greater emphasis on the likelihood of bad outcomes, thus skewing the perceived probabilities toward worse scenarios.

In the context of growth uncertainty, the pessimism effect shifts the focus toward worse technological outcomes, increasing the weight on scenarios where economic growth is slower and future welfare is lower. If we were to keep aversion to intertemporal substitution at $\eta = 2$, as in Nordhaus' original DICE scenario discussed above, low technology realizations lead to a higher marginal value of climate capital. In particular, Jensen and Traeger (2014) show that the increase in the marginal value of climate capital is more than three times greater than the increase in the marginal value of produced capital. As a result, the overall opportunity cost of abatement decreases, and uncertainty drives a significant increase in the optimal level of abatement and, correspondingly, the SCC. We reproduce the corresponding result in Figure 3's blue lines (with and without shock persistence) where $\eta = 2$ and $\text{RRA} = 10$.

Yet, the standard calibration of Epstein-Zin-Weil preferences suggestest a quite different value of the intertemporal elasticity of substitution. Once we disentangle it from risk aversion, the long run risk literature in asset pricing finds that aversion to mere intertemporal fluctuations is closer to $\eta = \frac{2}{3}$ than to $\eta = 2$. In contrast, relative risk aversion in these estimates becomes much larger and closer to values of $\text{RRA} \in [6, 10]$. In these scenarios, as the result of the low η, the pessimism term reduces the expected value of abatement, while simultaneously increasing the expected value of produced capital. This dynamic causes the expected opportunity cost of abatement to rise, and the pessimism weighting reduces the optimal level of abatement in this scenario. Here, the decision maker is more focused on preserving produced capital rather than investing in abatement, given that the relative value of climate capital is lower in these pessimistic states. We depict this result in the left panel of Figure 4. The right panel shows how the SCC and, thus, the risk premium changes as a function of the aversion to intertemporal substitution under Epstein-Zin-Weil preferences.

Figure 4 also shows that, despite the negative risk premium, the SCC increases w.r.t. to the original $\eta = 2$ scenario. As pointed out by Crost and Traeger (2014), this is a recalibration effect resulting from the lower growth discounting under Epstein-Zin-Weil preferences. These preferences split the aversion embedded in the unique utility function of the standard additive expected utility model into a lower aversion to intertemporal substitution and a higher aversion to risk. In a scenario with expected growth, aversion to intertemporal substitution increases the marginal utility of the present generation who have to pay for the abatement, relative to the marginal utility of the future generations who benefit from the abatement. Reducing η, therefore, reduces the valuation of the cost and increases the valuation of the benefits of abatement (reducing the growth discounting term of the consumption discount rate). Thus, Epstein-Zin-Weil preferences increase the SCC despite the negative risk premium they imply in the context of growth uncertainty. Epstein et al. (2014) point out that the common calibration of Epstein-Zin-Weil preferences can cause a rather high willingness to pay for an early resolution of uncertainty (even in the absence of any instrumental value). Traeger (2019) explains that the standard calibration and $\eta \leq 1$ can be explained by an artefact of isoelastic Epstein-Zin-Weil preferences, which behave somewhat counterintuitively for $\eta > 1$. This reasoning makes a choice of unity for the intertemporal elasticity of substitution under Epstein-Zin-Weil preferences somewhat attractive, in combination with a substantially higher coefficient or relative risk aversion. In the next section, we discuss how to solve the climate-economy model presented in Section 2.1 in closed form for this particularly attractive version of Epstein-Zin-Weil preferences. We note that, under growth uncertainty, Epstein-Zin-Weil preferences also imply a precautionary savings effect and refer to Jensen and Traeger (2014) for details.

3.7 Closed-Form Solutions

Hoel and Karp (2001, 2002) inspired what we might call the first generation of dynamic climate-economic models that analyze uncertainty using closed-form solutions. These "linear-quadratic" models feature a quadratic objective function and exclude an explicit production function. The climate system is simplified to a model of atmospheric carbon, although this limitation had been mitigated by the introduction of the TRCE model. These authors, along with many who followed, primarily used the framework to study issues related to Weitzman's (1974)'s analysis of prices versus quantities in a dynamic climate change context, concluding that in most cases, taxes are more effective than quotas.

In a highly stylized dynamic analytic climate-economy model, Weitzman (2009) derives a dismal theorem, effectively stating that we should dedicate all of our resources to combating climate change. As briefly discussed in Section 3.3 in connection to his earlier papers, this reasoning largely relies on a mitigation project being the only option to deliver a certain unit of consumption into a hugely uncertain future. Unfortunately, his "fait-tails" paper uses this interesting argument to discredit cost-benefit analysis in the climate context and to argue that uncertainty is more important than time preference when evaluating the climate problem. This reasoning has caused a lot of confusion in the field, so we briefly elaborate on the core intuition of the model, even if the model is far more stylized than the other models that we are discussing here. First, Weitzman starts out by describing the uncertainty over the future by an improper prior, whose expectations are infinity by assumption. Such priors can be helpful to analyze where a Bayesian parameter estimate bunches up its mass. Yet, by assumption, they defy an expected utility representation of preferences.[22] We could find alternative preference representations, such as lexicographic preferences that give immediate priority to one particular outcome. We could also interpret Weitzman's infinity as such a priority given to climate change. Yet, his model's priority really only applies to passing the first safe consumption unit into the future. Afterwards, we can resume standard cost-benefit analysis. Alternatively, if we believe that our doomsday climate scenarios make it impossible to pass a safe consumption unit into the future, the value of the mitigation project merely changes the uncertainty and might well be finite.

The decision-theoretic literature offers a lot of maxi-min type representations where we are very focused on a particular doomsday scenario. Yet, keeping in mind that climate change is not the only challenge humanity is facing on this planet, we are somewhat skeptical of how useful it is to abandon the type of trade-offs more carefully structured and normalized models allow us to do. In general, such trade-offs can be well-represented by cost-benefit analysis, which, of course, has to include proper treatments of uncertainty. Section 4.4 surveys models that pay attention to maxi-min-related ideas, including model misspecification and ambiguity aversion, and still deliver well-defined trade-offs. We interpret Weitzman's work instead as a call to better incorporate uncertainty analysis into cost-benefit analysis (rather than defying cost-benefit analysis). Costello et al. (2010) show that placing an upper bound on the temperature distribution renders a well-behaved model that is not particularly sensitive to the temperature cut-off. We refer to Millner (2013) for interesting variations of Weitzman's dismal theorem and a somewhat different critique. We further discuss learning under fat-tails in the context of Kelly and Tan's (2015) work in Section 4.2, and we re-examine the conjecture that uncertainty dominates time preference further below in a more structured model.

van den Bremer and van der Ploeg (2021) take the discussion of the previous sections into the continuous time framework and introduce a new and promising solution technique, common in the sciences, into climate change economics. The starting point is a simple AK-model with a known closed-form solution. This one state model has no climate but an analytic solution even in the case of their adopted Epstein-Zin-Weil preferences. The authors then introduce climate as a first order perturbation to the system. Given climate change damages are generally small, such a first-order perturbation gives a closed-form solution that can be conjectured a good approximation of the underlying AK-climate economy. The nice feature is that this first-order perturbation is itself a nonlinear function of the climate. While the abatement policy is exogenous, the authors can derive a complex parametric SCC formula discussing several of the phenomena of the previous sections jointly in parametric form (some as part of the discount rate and some as part of the price of carbon). In addition, they introduce exogenous correlations between different stochastic processes that give rise to a set of 'climate betas'. The authors find that growth uncertainty dominates the SCC's risk premium, which is positive in their base runs assuming $\eta > 1$. They also point out that growth uncertainty turns their SCC risk premium negative when calibrating η accordingly to the asset pricing literature ($\eta = \frac{2}{3}$).

[22] Most axiomatizations of expected utility theory assume in addition to the von Neumann-Morgenstern axioms that utility is bounded and that probabilities are properly normalized and integrate to unity. Both of these assumptions are violated in Weitzman's base model. We can permit unbounded utility at the expense of restricting how utility/consumption is permitted to change over time.

Golosov et al. (2014) sparked a new wave of interest in closed-form IAMs. Under an assumption of full capital depreciation and logarithmic utility, the authors show how to solve the optimal energy allocation problem in an economy with explicit production and a CES-substitutability between coal, oil, and renewable. For this purpose, they simplify the climate system to equation (3) governing the carbon cycle and argue that the strong concavity of the radiative forcing equation (4) cancels the convexity of damages in temperature. What this reasoning overlooks is that warming from atmospheric carbon happens with a substantial delay, which actually reduces the SCC, see below.[23]

Golosov et al.'s (2014) model has been criticized for ignoring capital in energy production, arguably one of the most important inputs in fossil and renewable energy production. Traeger (2015,2023) shows that the model can also be solved with the general production function from equation (3), which includes capital in energy production and can also vary energy substitutability by sector (Traeger, 2022). This Analytic Climate Economy (ACE) also relaxes Golosov et al.'s (2014) full depreciation assumption, which would replace the capital accumulation equation (7) by $K_{t+1} = I_t$. Such a full depreciation assumption has two main shortcomings: (i) it misses the remaining capital from previous periods, and (ii) it fails to account for how current investments increase future capital availability. ACE uses instead the equation

$$K_{t+1} = I_t \left[\frac{1 + g_{k,t}}{\delta_k + g_{k,t}} \right] \tag{19}$$

provides a "quick fix" using a correction factor that relies on an exogenous approximation of the growth rate of capital $g_{k,t} = \frac{K_{t+1}}{K_t} - 1$ and falls in the rate of capital depreciation δ_k. Traeger shows that equation (19) coincides with the standard capital accumulation equation (7) if the exogenous capital growth approximation is accurate (or if $\delta_k = 1$).[24]

In the context of uncertainty, Golosov et al. (2014) employ their analytic framework to evaluate the impact of a damage coefficient that, in some future year, will either jump up or down to a given value reflecting high or low damages. Motivated by the reasoning discussed in the preceding section, Traeger (2015) integrates Epstein-Zin-Weil preferences with an arbitrary degree of risk aversion into the framework, while keeping aversion to intertemporal substitution equal to unity. It is this difference between aversion to risk and to intertemporal substitution that is exploited by the asset pricing literature to explain observed risk premia. Moreover, Traeger (2015) points out that Golosov et al.'s (2014) damage function over CO_2 is slightly concave, making it less well suited for an analysis of uncertainty given that the damage function's convexity in temperature plays an important role as we saw in Section 3.5. Thus, he shows how to solve the IAM with a more comprehensive climate model following the standard equations (3-5) explained in Section 2.1. While it employs a standard carbon cycle or impulse response model for the carbon cycle, Traeger shows that an analytic solution requires a non-linear model for temperature dynamics, whose impulse response compares nicely with other models and CMIP 5 (see Figure 2).

Using the damage function presented in equation (2), Traeger (2015,2021) shows how to solve climate-economic models of the form presented in Section 2.1 under the assumptions above if the underlying

[23]There has been some confusion on this point resulting, likely, from equation (6)'s TCRE model. Cumulative historic emissions can be a good approximation for warming (under certainty), and the equation suggests that there is no warming delay. However, the approximation merely cancels the delayed warming resulting from *atmospheric* carbon's radiative forcing in the temperature equation with the slow removal of carbon from the atmosphere in the equation characterizing carbon dynamics.

[24]We note that most IAMs calibrate an exogenous rate of technological progress incorporated into $\boldsymbol{A}_t$ to observed growth, making this exogeneity less of an issue than one might think. Yet, one has to be careful when calibrating the model for short time steps.

uncertainty is governed by affine stochastic processes. Affine stochastic processes belong to a class of stochastic processes that have a particular structure: an additive decomposition of the cumulant generating function. Cumulants are close relatives to the standard central moments of a probability distribution. The first cumulant is the expected value, the second the variance, the third captures skewness, the fourth kurtosis, and so on. Examples of affine stochastic processes include autoregressive processes with somewhat arbitrary shock distributions or a normal-normal learning model. Affine processes have been widely used in asset pricing models due to their analytical tractability.

Lemma 1 in Traeger (2021) transforms Section 2's recursive stochastic dynamic programming problem on the state space into a system of (nonlinear) algebraic equations for the shadow values, making the numeric problem significantly easier to solve. For some specifications, these algebraic equations can be solved in closed form, providing valuable analytic insights into the impact of uncertainty. Relating back to Weitzman (2009), the paper shows that a simple additive autoregressive shock model describing carbon flow uncertainty leads to the following expression for welfare loss:

$$\Delta W \sim \sum_{l=1}^{\infty} \frac{1}{l!} \kappa_l \left(\frac{\beta}{1-\gamma\beta} \alpha \varphi_{M,1} \right)^l ,$$

where γ is the autoregressive coefficient, which measures the persistence of shocks over time, β is the discount factor, and $\alpha\varphi_{M,1}$ is the (intertemporal) risk-aversion weighted shadow value of atmospheric carbon under certainty. In this equation, κ_l represents the cumulants ($\approx$ moments) of the underlying shock distribution. First, we observe how the persistence of the shocks increases the welfare loss. This equation emphasizes that, in the context of climate change, we are worried mostly about long-term growth uncertainty and less about business cycles. In contrast to persistent shocks, iid shocks to emissions and carbon flows are smoothed out over the long time horizons relevant for climate change. Shock persistence derives naturally in the case of epistemological uncertainty, i.e., those situations where uncertainty is merely reflecting our ignorance of the true system dynamics rather than nature's stochasticity (Traeger, 2015). Then, observing a high realization in some period, also makes it more likely to observe a high realization in future periods.

Second, we observe the close interaction between time preference and the order l of the moment of the distribution. For autoregressive shocks, $\left(\frac{\beta}{1-\gamma\beta}\right)^l$ can be large and grows with l. Thus, time preference plays a particularly crucial role in the context of higher moments and the contribution from "fat-tails". Unfortunately, uncertainty does not make our models less sensitive to the heated debate about the right calibration of time preference, but rather more sensitive.

Turning towards the impact of uncertainty on the SCC, Traeger (2021) models three different stochastic processes governing temperature, carbon dynamics, and damages. Each of these processes reflect epistemological uncertainty implying autoregression of the shocks and an endogenous evolution of the uncertainty. For example, the stochastic carbon dynamics can be reasonably well approximated with a common stochastic process in the long-run risk literature on asset pricing. For temperature uncertainty, the model requires an autoregressive gamma process to match the climate response adequately.[25] In the

[25] Interestingly, the main reason is not so much the long-run skewed equilibrium climate sensitivity, but the non-linearities of the temperature system and the necessity to match expected temperature dynamics to those of the deterministic model for a fair evaluation of the risk premium.

absence of uncertainty, the model's deterministic social cost of carbon is

$$SCC_t^{\text{det}} = \frac{\beta Y_t^{\text{net}}}{M_{\text{pre}}} \xi_0 \underbrace{\left[(\mathbf{1} - \beta\boldsymbol{\sigma})^{-1}\right]_{1,1} \sigma^{\text{forc}}}_{\text{temperature dynamics}} \underbrace{\left[(\mathbf{1} - \beta\boldsymbol{\Phi})^{-1}\right]_{1,1}}_{\text{carbon dynamics}}, \tag{20}$$

where $[\cdot]_{1,1}$ refers to the first element of the inverted matrix. The formula depends on several key factors. Its base is the discounted (net of climate damages) output of the economy Y_t^{net} multiplied by the damage (semi-)elastiticy ξ_0 and divided by the the preindustrial carbon concentration M_{pre}. The part $\sigma^{\text{forc}}(\mathbf{1} - \beta\boldsymbol{\sigma})^{-1}$ captures the temperature dynamics and the term $(\mathbf{1} - \beta\boldsymbol{\Phi})^{-1}$ describe the persistence of carbon dynamics, respectively.[26] The temperature dynamics multiplier reduces the optimal carbon tax by approximately 20-40% because ocean cooling provides a delay in warming, whereas the carbon dynamics multiplier increases the carbon tax by a factor of 4 and more depending on the rate of time preference as a result of long life-time of atmospheric carbon (Traeger, 2023).

Here, we are interested in how uncertainty increases this SCC. Let $\theta_i(z_i)$ denote the percentage increase of the SCC as a result of adding uncertainty $i \in T, M, d$ to either temperature dynamics, carbon flow dynamics, or damage and damage-adaptation dynamics. Then, adding an individual uncertainty to the model, we have

$$SCC^{\text{unc}} = SCC^{\text{det}} \left(1 + \theta_i(z_i)\right).$$

We will not reproduce the detailed formulas for each uncertainty here, but summarize the main findings and use the model to discuss the interaction of multiple uncertainties.

The relative risk premium $\theta_T(z_T)$ from temperature uncertainty is a generally increasing and convex function of the risk aversion weighted shadow value of carbon under certainty. Thus, the *relative* importance of risk is larger, the larger the level of the SCC. As we discussed for welfare above, also the SCC increases strongly with the shock persistence and discount factor (= reduction of the rate of pure time preference). Pure time preference is far more important than the general discount rate because growth discounting is cancelled by a proportional damage increase. Moreover, the SCC's risk premium from temperature uncertainty can be divided into two contributions. The first stems directly from the convexity of the damage function in temperature, and the second interacts these damages with the type of risk aversion encapsulated in the Epstein-Zin-Weil preferences (intertemporal risk aversion).

The relative risk premium $\theta_M(z_M)$ from carbon flow uncertainty is also increasing and convex in the risk aversion weighted shadow value of carbon under certainty, with the same implication that the SCC's relative risk premium grows convexly in the level of SCC^{det}. However, in contrast to the risk premium from temperature uncertainty, for carbon flow shocks, the damage function's convexity channel is mostly canceled by the strong concavity in equation (4), which translates CO_2 emissions into radiative forcing. As a result, carbon flow uncertainty has a somewhat negligible impact on the SCC, a finding emphasized by Traeger (2023) as follows. The persistence of CO_2 in the atmosphere implies that, under certainty, the main multiplier of the SCC in equation (20) is the carbon dynamics multiplier whereas temperature's warming delay reduces the SCC. By contrast, under uncertainty, the concavity of the

[26] As explained in Traeger (2023), a Neumann series expansion of $(1 - \beta\boldsymbol{\Phi})^{-1} = \sum_{i=0}^{\infty} \beta^i \boldsymbol{\Phi}^i$ shows that the expression in square brackets characterizes the discounted sum of carbon that remains in or returns to the atmosphere over an infinite time horizon while cycling through the different reservoirs, and similarly for the heat exchange matrix $\boldsymbol{\sigma}$

radiative forcing equation implies that carbon flow uncertainty delivers an almost negligible risk premium whereas temperature uncertainty contributes substantially.

This finding on the relative importance of temperature uncertainty and the relative unimportance of carbon flow uncertainty explains why we decided to switch away from the TRCE model when we started discussing climate uncertainty in Section 3.5. The TRCE model gives us a convenient approximation for the temperature response to emissions. However, it masks the different non-linearities that play a crucial role for understanding the impacts of different uncertainties in the climate system.

Finally, damage uncertainty results in a related relative risk premium θ_d. We want to focus now on the interaction of these three different climate uncertainties. Incorporating the implications of these three uncertainties jointly into the model delivers

$$SCC^{unc} = SCC^{det} \cdot \Big(1+\theta_d\Big) \cdot \Big(1+\theta_T\big(z_T \cdot (1+\theta_d)\big)\Big) \cdot \bigg(1+\theta_M\Big(z_M \cdot \Big(1+\theta_T\big(z_T \cdot (1+\theta_d)\big)\Big)\Big)\bigg) \quad (21)$$

While each uncertainty type independently raises the SCC, their interaction causes an even larger increase. First, uncertain damages also increase the impact of an uncertain temperature as seen in the new relative temperature risk premium $\theta_T\big(z_T \cdot (1+\theta_d)\big)$ whose argument is amplified directly by the damage uncertainty multiplier $(1+\theta_d)$. Similarly, temperature and damage uncertainty both amplify the social cost of carbon fluctuations, leading to the final multiplier in equation (21). Recalling that the functions $\theta_T(\cdot)$ and $\theta_M(\cdot)$ are generally increasing and convex in their arguments, these relative risk premia grow in the case of joint uncertainties. In addition, the mere fact that these relative risk premia interact multiplicative implies another mutual aggravation of the uncertainty premia, which is also convex in the sense that $(1+x\%)^3$ is marginally larger than $3x$ for small x, but substantially larger than the sum of the individual contributions if each contribution by itself is already large. Traeger finds that only the temperature premium is individually sizable. The mutual aggravation of the premia is noticeable, but only substantial in the case of low rates of pure time preference.

4 Numeric Modeling and Insights

Most quantitative studies use numeric methods as the set of models with closed-form solutions is limited. The majority of studies exploring uncertainty started out with scenario analysis that eventually was expanded into what we label "Monte-Carlo simulations" discussed in Section 4.1. Incorporating optimization and forward-looking behavior generally requires numeric dynamic programming, an approach that we will discuss together with some of the literature's insights on anticipated learning in Section 4.2. Thus far, we have discussed smooth and reversible stochastic events. Section 4.3 on tipping points discusses the modeling, literature, and results governing large-scale irreversible climate phenomena that can be triggered by global warming. Section 4.4 takes a closer look at a literature where modelers have been weary to assume given probabilities or unique models to analyze climate economics. Finally, Section 4.5 gives a very brief introduction to a recent approach that harnesses the power of machine learning to solve more complex dynamic programming problems.

4.1 Monte-Carlo Simulations

Monte Carlo simulations are the simplest and, at times, the only feasible way to introduce uncertainty into complex policy-oriented IAMs. These studies draw a large number of separate worlds and compare results statistically. The simulations generally assume some given policy regime and simulate the IAM forward. They introduce uncertainty by drawing realization from a joint probability distribution representing uncertainty about multiple model parameters. Repeated drawing and simulation generate distributions of how the future might look under a given policy. In these models, the SCC is commonly calculated by running the model twice for each draw of the parameter distribution: once for the given policy scenario and a second time with an additional 'pulse' of CO_2 emissions. The SCC is then determined by comparing the economic loss generated by an additional ton of carbon. Like for all outcome variables, Monte Carlo simulations also generate a distribution over the SCC. Its average value is often interpreted as the SCC under uncertainty.

The strength of the Monte Carlo approach is that it is relatively simple to implement and permits the introduction of multiple uncertainties into complex IAMs using realistic climate modules and detailed, disaggregated representations of the world economy. Such complexity is challenging to handle with stochastic dynamic programming approaches. The results explore possible futures and indicate the relative importance of different uncertainties and the overall sensitivity to parameter changes.

The downside of the Monte Carlo approach is that it cannot optimize policies. In a similar vein, future policymakers do not respond to the resolution of uncertainty, which can introduce a sort of 'negativity bias' into the simulated futures. This contrasts sharply with Bellman's approach characterized by equation (10) in section 2.2, where the policymaker optimizes in every period conditional on the given information set (we might well consider this the optimistic extreme). Sometimes, modelers try to "fix" the lack of optimization by simply optimizing the policy of each run by itself. As shown by Crost and Traeger (2013), this approach is generally a bad idea. In their example of damage uncertainty, it can result in the wrong sign of the uncertainty effect. The issue is that, in such an approach, each decision maker lives in her own separate world, knows her particular future with certainty, and optimizes accordingly, which is quite different from optimizing in the face of uncertainty where decisions have to trade off the consequences across all possible futures.[27]

The trouble with deriving a distribution of SCCs is that setting today's policy requires a unique policy recommendation. As we know from section 3, the optimal response to uncertainty is not equal to the average response. It is based on optimizing expectations, as done by the Bellman equation (10). This always gives rise to a unique SCC conditional on current information.

Already the Stern Review (Stern, 2007) includes Monte Carlo simulations carried out with the PAGE IAM (Hope, 2006). Nordhaus (2008) uses this approach for his DICE model. Using the baseline scenario (no emission reduction policy), he first draws randomly 100 times from the iid normal distributions of eight parameters[28] separately. He simulates the (slightly simplified) DICE model for each set of realizations. In this exercise, he finds that the quantitatively most important variables are productivity growth and

[27] As we discussed in Section 3.4, it can also happen with standard (non-optimizing) Monte Carlo analysis that the sign differs between a world where policies are optimized and respond to uncertainty and where we merely simulate an expected SCC along a fixed policy trajectory. Yet, the issue with (simplistically) optimizing policy is that decisions can only be based on the current information set (see Footnote 7), which is violated when letting each decision maker foresee its actual future.

[28] The eight parameters are: growth in factor productivity, rate of decarbonization, equilibrium climate sensitivity, damage coefficient, backstop technology price, asymptotic population, transfer coefficient in carbon cycle and total fossil fuel resources.

climate sensitivity - a one standard deviation from their mean value has the largest impact on outcomes. He then repeats the exercise with 100 runs of the parameters being jointly uncertain to find the overall impact of the eight uncertainties and analyze potential interactions. Again productivity uncertainty dominates in his analysis and leads to a negative risk premium.

Gillingham et al. (2018) run and compare Monte-Carlo simulations for six common IAMs[29] for uncertainty about productivity growth, population growth, and the equilibrium climate sensitivity. Using 125 runs for each model, they reproduce Nordhaus's earlier result that productivity growth has the largest impact of the three parameters. They also conclude that the parametric uncertainty is more important than the differences between the six models, which they interpret as model uncertainty.

Most prominently, the Monte Carlo method has been used to calculate distributions, percentiles, and a mean SCC for US government agencies to evaluate the emission impact of large projects and new regulations in governmental cost-benefit analysis. In 2010, 2013, 2016 and 2021, the Interagency Working Group on the Social Cost of Carbon[30] simulated three IAMs (DICE, FUND, and PAGE) to create distributions of the social cost of carbon for three different discount rates and five different socio-economic scenarios with 10,000 draws from a probability distribution for the climate sensitivity.[31] The latest official set of numbers can be found in EPA (2023), which is based on input from the following two studies.

Rennert et al. (2022) introduce a new IAM, GIVE, with the explicit purpose of updating the SCC (and other SC-GHGs) used by US government agencies. The GIVE model improves earlier models along several dimensions. First, building on Rennert et al. (2021), they substantially improve the representation of economic uncertainty. Rennert et al. (2021) combine recent empirical work with expert elicitation to generate long-term probabilistic projections of population dynamics, economic growth, and emissions. Additionally, rather than drawing parameters once before a run, these simulations generate 'stochastic paths' for some of the uncertain variables. In addition to the probabilistic projections, the 10,000 Monte Carlo draws in Rennert et al. (2022) also incorporate a wide range of climatic and damage-related uncertainties. Second, the authors develop bottom-up country-level damage functions based on recent empirical work on damages from agriculture, mortality, sea-level rise, and energy. Third, they use the reduced complexity climate model FAIR as their climate module.

In parallel, the researchers of the The Climate Impact Lab (2023) developed their own high-resolution IAM, DSCIM. Their model also uses the probabilistic projections by Rennert et al. (2021) and the FAIR climate model, deploying Monte Carlo simulations similarly as outlined above. The Climate Impact Lab has made outstanding econometric efforts to empirically estimate the impacts of climate change, with the aim of putting the damage function on the best empirical foundations possible. For (so far) five impacts (health, and mortality, agriculture, labor, energy, and coastal impacts), they use highly granular temperature and outcome data from around the globe and state-of-the-art econometric methods to estimate and simulate the damages (Burke et al., 2015; Hsiang et al., 2017; Rode et al., 2021; Carleton and Greenstone, 2022; Carleton et al., 2022). All of these studies demonstrate the large uncertainties surrounding climate change and the SCC.

[29] DICE (Nordhaus and Sztorc, 2013), FUND (Anthoff and Tol, 2013), GCAM (Calvin et al., 2011), MERGE (Manne et al., 1995), IGSM (Webster et al., 2012), and WITCH (Bosetti et al., 2006).

[30] Interagency Working Group on Social Cost of Carbon (2010, 2013, 2016); Interagency Working Group on Social Cost Greenhouse Gases (2021).

[31] For PAGE and FUND, additional model parameters are uncertain in those simulations.

4.2 Classic Numeric Stochastic Dynamic Programming and Learning

In recent years, numeric stochastic dynamic programming has become the most widely adopted approach to analyze optimal climate policies under uncertainty and learning. This approach's popularity is due to two main reasons. First, only a limited number of problems can be solved analytically in closed form (Section 3.7). Second, the simpler Monte Carlo method described in Section 4.1) does not optimize policy or permit future policymakers to respond to new information. As a result, they serve better for stochastic scenario exploration than to pin down the optimal policy in the face of uncertainty. A seminal paper by Kelly and Kolstad (1999) translated the DICE model into a recursive numeric stochastic dynamic programming model already 25 years ago. It took more than a decade until their pioneering approach became more widely adopted in the field of climate change economics.

The crucial challenge of dynamic programming is finding the value function. Numerically, we approximate the value function using a set of basis functions, which represent the value function as a linear combination of simpler, predefined functions. These basis functions are often polynomials or splines. The value function is then expressed as $V(\boldsymbol{X}) \approx \sum_{i=1}^{n} \theta_i \phi_i(\boldsymbol{X}) \equiv V_\theta(\boldsymbol{X})$, where θ_i are coefficients to be determined and $\phi_i(\boldsymbol{X})$ are the basis functions. Using the approximation of the value function, the Bellman equation from Section 2.2 reads

$$V_{t,\theta}(\boldsymbol{X}_t) = \max_{\boldsymbol{a}_t} \; u(C_t) + \beta \; \mathbf{E}_t \left[V_{t+1,\theta}(\boldsymbol{X}_{t+1}) \mid \boldsymbol{X}_t \right], \tag{22}$$

subject to the climate-economic equations of Section 2.1. Here, we made the informational states $\boldsymbol{I}_t$ part of the state vector $\boldsymbol{X}_t$ for notation convenience.

With a finite time horizon, the value function becomes period-dependent and picks up a time index, $V_{t,\theta}$, which we already added to equation (22). In the final period T, the value function is simply the utility from consuming whatever is left/available based on the system state $\boldsymbol{X}_T$. We can then iterate backwards using equation (22). For this purpose, we have to define at which points to evaluate the value function. Unfortunately, numeric evaluation at all points is computationally infeasible, and we do not yet know where the optimally controlled system will arrive in the final period. Therefore, we must define a reasonably large state space that should cover the relevant possible states, e.g., based on some (stochastic) steady state reasoning or by starting with a large state space and refining it later. On this state space, we have to decide on a grid of points at which we evaluate the right side of equation (22), called the interpolation nodes. Evaluation at these grid points should give us a somewhat dense evaluation of the right side of equation (22), sufficiently so to get enough evaluation points allowing us to then fit the previous period's value function $V_{t,\theta}(\boldsymbol{X}_t)$ with reasonable precision. We obtain the basis coefficients of θ by some form of projection method, e.g., least squares, or if the number of coefficients equals that of the gridpoints, we obtain them from a simple matrix inversion (called collocation). This way, we can work our way backwards to the present and, in principle, are done. In practice, we want to carefully check the quality of our solution using a variety of methods going beyond this simple chapter.

The crucial step we skipped above is how to evaluate the uncertainty. Here, we can use a Monte Carlo method, merely drawing a set of realizations and calculating expectations. More computationally efficient is usually the approximation of the uncertainty distribution by so-called Gauss-Legendre nodes that satisfy certain properties, ensuring a good approximation of the distributional moments with as few evaluations as possible. Because we only have to calculate ahead one period at a time, the uncertainty evaluation is simple, and solving the stochastic problem is hardly more difficult than solving a deterministic problem.

With an infinite time horizon, we do not require the time index on the value functions because all periods look the same. If some equations of motion are non-stationary, we can always include the state representing the underlying change explicitly (which could be population, the technology level, or even time itself). Finding an approximation to the value function is still similar to the case of a finite time horizon. Instead of a last period, we start with a guess of the value function, which could simply be zero. Then, as before, we choose (i) the intervals over which to approximate the value function (relevant state space), (ii) the interpolation nodes at which we evaluate the optimization problem in the Bellman equation, and (iii) the basis functions with which to approximate the value function.[32] We then iterate the Bellman equation until we (hopefully) converge to a solution where the right side equals the left side of the Bellman equation up to a small approximation error.

The convergence of this solution approach is guaranteed in theory, conditional on a well-behaved model that implies a contraction mapping. However, the theoretical convergence results assume that we evaluate the system everywhere in our state space and do not approximate the value function either.[33] Yet, in practical terms, convergence might fail. We only approximate the function with a limited number of basis functions and only over a subset of the state space. In particular, potential 'kinks' in the value function can be hard to approximate, and sometimes transitions lead us out of the approximated state space where our value function evaluation becomes quite error-prone. The upside of value function iteration in the infinite time horizon setting is that if the algorithm converges, we generally have a reasonable solution conditional on checking some robustness criteria.

The main drawback of (classical) numeric stochastic dynamic programming is the curse of dimensionality, limiting how detailed a climate-economy model we can represent. To capture climate and economic dynamics faithfully requires several state variables. In addition, uncertainty and learning require informational states. In the standard approach, increasing the number of state variables increases the computational burden exponentially. Suppose we have three ($n = 3$) state variables: capital, carbon stock, and time (to capture exogenous processes). Suppose we use a fifth-order Chebyshev polynomial in each dimension, including all interaction terms, and, accordingly, use five interpolation nodes along each dimension ($m = 5$) to approximate the value function. Then, for each iteration, we must approximate the value function at $m^n = 125$ grid points and optimize the Bellman equation 125 times. Adding states with the same polynomial order leads quickly to prohibitively large computations. For our example, doubling the number of states to six requires $15,625$ nodes, a 125-fold increase in nodes. Adding another three states would require $1,953,125$ nodes, and so on. A similar increase in the computational burden happens if we want to improve the precision of the approximation by adding to the order of the Chebyshev polynomials for one or more states.[34] There are a variety of methods to reduce the computational burden of the dynamic programming problem. For example, using Smolyak grids or endogenous grids reduces the computational burden to only grow polynomially in the state space, and Howard's method suggests that we don't have to reoptimize the right side of the Bellman equation in every period. We also note that dynamic programming is particularly well-suited for parallelization as the optimization problem can

[32] We generally choose the basis functions and interpolation nodes jointly. E.g., Chebyshev polynomials should better be evaluated at the Chebyshev nodes. Combining Chebyshev polynomials with equidistant nodes might give rise to Runge's phenomenon: large approximation errors at the boundaries of the state space

[33] Stachurski (2009) gives sufficient conditions for convergence under a linear approximation of the value function.

[34] For six states, adding one node in one state space dimension adds 3125 total grid points; increasing by one node in each state space dimension requires a total of $46,656$ grid points. Several approaches to circumvent the constraints imposed by computational power exist, and we discuss them and their characteristics in section 4.5.

be easily spread out over different computation nodes. We refer the reader to Judd (1996) and Stachurski (2009) for more details and briefly summarize one of the most promising recent approaches based on machine learning techniques in Section 4.5.

Kelly and Kolstad (1999) and Kelly and Kolstad (2001) use neural nets to approximate the value function to solve the stochastic DICE model in 24h. For using a Sun Spark 20 with 75mHz at the time, they use an impressive grid of 57,600 points, and, yet, the policy function approximation using the neural network is a bit coarse, and even the sign of the optimal policy response is not quite obvious. Traeger (2012) finds that a similar DICE-based model with the same amount of state variables as in Kelly and Kolstad's (1999) original implementation can be solved to high precision using fewer grid points employing either Chebychev polynomials or splines for the value function approximation. Given the smoothness of the solution, Chebychev polynomials performed the best.

Starting with Kelly and Kolstad (1999), the stochastic DICE model has been recognized as a helpful framework to investigate the dynamics of learning over model parameters and, in particular, the equilibrium climate sensitivity. These models commonly assume that the optimizing decision maker employs Bayesian updating. The decision maker holds an initial belief (a prior) governing the climate sensitivity and updates this belief using Bayes' Theorem (generating a posterior or updated prior). The prior is a subjective probability distribution over all the possible values of the equilibrium climate sensitivity. If this original prior is (or can be approximated by) a so-called conjugate prior, the updated prior has the same distribution as the original one. Then, we can fully describe the uncertainty using this distribution's parameters. For a normally distributed prior, for example, the epistemological uncertainty is fully captured by two informational state variables, the prior's mean and variance. The updating process depends on how new information about the parameter arrives. Suppose the decision maker observes stochastic temperature realizations and knows the structure of the temperature equation and the distribution of temperature shocks (likelihood function). Then, with every temperature observation, she can infer an estimate of the 'realized' equilibrium climate sensitivity and update her prior.

Kelly and Kolstad (1999) investigate how long it takes to learn the 'true' value of equilibrium climate sensitivity, which is about 90 years in their framework, and for their definition of the 'true' value. They point out that learning speed positively (if weakly) depends on emissions. Given limited numeric precision, the authors discuss learning more than the optimal uncertainty response of the policy, which in their plots would likely suggest a negative risk premium from climate uncertainty under learning. Karp and Zhang (2006) investigate learning in the linear-quadratic model with a mostly analytic solution. They find that anticipated learning in their setup decreases the optimal level of abatement relative to a situation with uncertainty only uncertainty. Leach (2007) introduces a second uncertain variable, the persistence of temperature changes, into the model by Kelly and Kolstad (1999) and shows that learning slows down significantly - by hundreds of years. His findings suggest that the decision maker increases emissions to accelerate learning. Jensen and Traeger (2013) revisit uncertainty about the climate sensitivity in the DICE model, using Traeger's (2012) insights on numeric precision. They find a positive risk premium and precautionary emission reductions and no emission response to the ability to learn (no 'wait-and-see').

Daniel et al. (2019) build a stylized model with Epstein-Zin-Weil preferences and uncertainty. The DICE-inspired framework features six time periods, six possible states of the climate, and technological progress. They find that the SCC declines over time as uncertainty declines, a finding they attribute to learning under Epstein-Zin-Weil preferences. We note that Jensen and Traeger (2013) also feature runs using Epstein-Zin-Weil preferences in a more comprehensive implementation of DICE with Bayesian learning and find the usual increase of the SCC over time.

Inspired by Weitzman (2009), Kelly and Tan (2015) introduce fat-tailed uncertainty over climate sensitivity into the DICE model. They find that learning the true value of climate sensitivity is also a comparatively slow process in their model, confirming results from Kelly and Kolstad (1999) for a normal climate sensitivity prior. However, they also find that Bayesian learning from observed temperatures leads the decision maker to quickly discard the fat tail of the distribution (if the underlying true climate sensitivity is not in the tail). They also investigate the impact of uncertainty *without* learning in a counterfactual model in which the decision maker does not update their prior despite new information. The risk optimal present-day SCC is higher under uncertainty and learning than under certainty (for the expected value of the climate sensitivity), and much higher still under uncertainty without learning. Like Kelly and Tan (2015), Hwang et al. (2017) analyze learning under fat-tailed risk. They use a somewhat different DICE-based model, a different solution method, and a different learning definition. They find that anticipated learning lowers the optimal SCC compared to a situation with uncertainty but no learning. They also find that the optimal policy is less sensitive to the uncertain variable's true value than in the pure uncertainty case. Their solution method is somewhat idiosyncratic. The authors assume that the value function is an additively separable combination of logarithms in each dimension, which mostly eliminates the usual computational burden.

In Rudik (2020), the decision maker is uncertain about the damage function coefficient and exponent. She does not learn about the coefficient but about the exponent. Immediate learning is hindered by stochastic multiplicative shocks to damages. The decision maker learns the exponent value over time from observing realized output (net of damages) and temperatures. Overall, the contributions of damage uncertainty and learning to the present-day optimal SCC are small but positive. Kelly et al. (2024) investigate uncertainty and learning about two parameters simultaneously: the effectiveness of solar geoengineering and the climate feedback parameter. Their main numeric finding is that at the optimal level of solar geoengineering deployment[35] uncertainty slows learning about solar geoengineering effectiveness and climate sensitivity. They explain this effect analytically. In principle, 'large' deployment of solar geoengineering can cause big temperature changes that speed up learning. However, the 'small' optimal deployment makes overall temperature changes smaller, and the joint uncertainty reduces the learning signal of the temperature observation. Meier and Traeger (2023) examine solar geoengineering in an extension of the ACE model that distinguishes long-run persistent uncertainty governing climate and some aspects of solar radiation management from quickly resolving components of the uncertainty governing the efficiency of and damage from solar geoengineering. They find that the quickly resolving uncertainty slightly lowers optimal initial deployment and that the conditional deployment after uncertainty varies widely depending on the realized learning, including large probabilities on no further deployment but also on quite high deployment.

[35]The optimal level is a 'small' deployment leading to a four percent reduction in radiative forcing (at expected effectiveness and climate sensitivity values). The paper relates closely to Heutel et al. (2018) who introduce solar geoengineering into a DICE-based model, and quantify the impacts of climate sensitivity uncertainty (increasing optimal deployment) and uncertain damages from solar geoengineering (decreasing optimal deployment). They use an approximate dynamic programming method to solve their model under uncertainty.

4.3 Tipping Point Models

In the realm of climate change, tipping points refer to situations where small changes in environmental conditions can lead to significant and often irreversible shifts in the Earth's climate system. Such regime shifts in the climate dynamics are generally triggered at an unknown threshold (if they occur at all). Unlike gradual change, these thresholds require IAMs to model discrete and irreversible changes in their system dynamics. This section explores the economic implications of tipping points, the role of uncertainty, and their policy implications.[36] Examples of tipping points include the melting of permafrost, releasing methane into the atmosphere; the release of methane hydrates from the ocean; an Amazon rainforest dieback, which releases stored CO2; and the melting of the Greenland ice sheet, which could raise sea levels by about 7 meters. The collapse of the Western Antarctic ice sheet could add about 2-3 meters to sea levels, while a weakening of the Atlantic Meridional Overturning Circulation, a crucial ocean current, could substantially alter weather patterns over Europe and around the globe. Climate change could also alter the Indian summer monsoon, whose timing is crucial for agricultural production, and alter carbon uptake by soils and the biosphere.

Uncertainty is particularly crucial in the context of modeling tipping points. If we were to know the tipping threshold, the optimal policy would simply be to stay underneath the threshold if worthwhile or cross it otherwise. While the question of whether or when to cross the threshold might still be of interest, we have not encountered any example in the climate system where the threshold is known. As a result, the task of managing tipping points becomes one of managing the hazard rate, i.e., managing the risk of tipping into a different climate regime. Following early literature on stylized tipping models (Heal, 1984; Clarke and Reed, 1994; Tsur and Zemel, 1996; Brock and Starrett, 2003), Keller et al. (2004) introduce a tipping point into the DICE model. Crossing a known threshold triggers a decline in the North Atlantic circulation impacting Northern Europe, which the authors represent as a damage whose magnitude is stochastic. Most importantly, they introduce uncertainty over the climate sensitivity implying that avoiding the threshold is a question of reducing the tipping hazard. The authors show that the optimal abatement rate increases in the presence of a tipping point. However, uncertainty might not lead to a larger reduction than if we could deterministically steer clear of the tipping point.

Keller et al. (2004) use the classic implementation of the deterministic DICE model and jointly optimize over an ensemble of DICE runs. Pathbreaking at the time, this approach significantly limits the informational structure of the decision maker and uncertainty cannot smoothly unfold over time. In their scenario where the decision maker can learn and respond to uncertainty resolution, he can only do so in the year 2085, learning that climate sensitivity takes on one of three different values. These shortcomings can be resolved using the dynamic programming approach of Sections 2.2 and 4.2, extending the approach to include tipping points. In the IAM context, this approach was first implemented by Lemoine and Traeger (2014) in a recursive implementation of the DICE model. Extending Kelly and Kolstad's (1999) pioneering implementation of the stochastic DICE model, the authors consider two possible regime shifts in the climate system. The first regime shift triggers temperature feedbacks that increase the climate sensitivity; the second regime shift reduces the carbon sinks. The authors model one of these tipping points at a time. The simple trick is to first solve the post-tipping regime the way described in Section

[36] In environmental economics, the careful modeling of tipping points likely began with the study of Skiba points, which emerged in the context of shallow lakes. This research explored how the sustained use of fertilizers can cause a lake to suddenly tip into a eutrophic state–a condition where the lake becomes oxygen-deprived, smelly, and largely devoid of life (e.g., Skiba, 1978; Brock and Starrett, 2003; Mäler et al., 2003; Wagener, 2003).

4.2, where post-tipping means that we are in a world where the irreversible change in climate system dynamics has already occurred. Let us call the resulting value function $V_1(\cdot)$. Then we move back to solve the optimal policy in the pre-tipping world using the modified Bellman equation

$$V_0\left(\boldsymbol{X}_t\right) = \max_{a_t} \Bigg\{ u\left(C_t\right) + \beta_t \mathbf{E}_t \Bigg[\left[1 - h\left(\boldsymbol{X}_t, \boldsymbol{X}_{t+1}\right)\right] \; V_0\left(\boldsymbol{X}_{t+1}\right) + h\left(\boldsymbol{X}_t, \boldsymbol{X}_{t+1}\right) \; V_1\left(\boldsymbol{X}_{t+1}\right) \Bigg] \Bigg\} \tag{23}$$

which we solve subject to the pre-tipping system dynamics. Neglecting the expectation operator for a moment, the hazard probability $h(\boldsymbol{X}_t, \boldsymbol{X}_{t+1})$ gives the (endogenous) probability of tipping. This probability depends on both the warming over the course of the period and past information capturing potential learning, captured by the states at time t and their change w.r.t. period $t+1$. E.g., Lemoine and Traeger (2014) model learning as an updating of a uniform Bayesian prior over the possible threshold locations, information that can be inferred from the current state of the physical system, avoiding the need for an additional (costly) state variable. Learning is important because it implies that, with some probability, we can avoid tipping by stabilizing temperatures at, say, $2°C$, a number promoted in the Paris Accord and closely based on the idea of trying to avoid dangerous tipping points.

The Bellman equation (23) states that our current value is composed of our current consumption utility plus the future value. The future value is now a combination of the hazard-weighted future value if we tip into the regime shift and the future value if we avoid tipping for another period as captured by the value function $V_0(\cdot)$. We stay in the pre-tipping regime precisely with the probability $1 - h\left(\boldsymbol{X}_t, \boldsymbol{X}_{t+1}\right)$. The expectation operator governs other non-tipping uncertainties and not the tipping hazard that is explicitly captured by $h(\cdot)$. In the examples of Lemoine and Traeger (2014) such uncertainty governs temperature stochasticity.[37]

The power of a recursive formulation of the tipping point problem is that it allows us to model the possibility of tipping in any possible period and to let the decision maker respond to a potential tipping or learn that she is in a safe region. It also allows us to combine tipping points with other stochastic events, such as temperature uncertainty, all of which would be unfeasible if we actually had to calculate through all the possible trajectories of the corresponding decision tree (see Figure 1).

Lemoine and Traeger (2014) distinguish between two drivers of the pre-tipping policy. First, the presence of a tipping hazard implies that today's optimal policy has to be tailored to its implications in both scenarios, tipping and remaining in the pre-tipping regime. As compared to a world without tipping, this reasoning adds a new term proportional to $DWI \equiv h\left[\frac{\partial V_0}{\partial \boldsymbol{X}_{t+1}} - \frac{\partial V_1}{\partial \boldsymbol{X}_{t+1}}\right] \frac{\partial \boldsymbol{X}_{t+1}(e_t^*)}{\partial e_t}$ to the SCC, which the authors label the *differential welfare impact (DWI).* An additional unit of emissions changes the future state of the system. With the hazard $h(\cdot)$ of tipping, we have to evaluate the effect of this change in the post rather than the pre-tipping regime. For example, if the tipping point suppresses the natural carbon sinks, then the welfare cost of increasing the CO_2 concentration today is worse in the post-tipping regime than in the current regime. Then, accounting for the differential welfare impact makes the decision maker more precautionary.

The second factor changing the SCC in the face of a possible tipping point is what Lemoine and Traeger (2014) call the *marginal hazard effect.* It accounts for the policy's implication on the hazard of

[37]The authors assume that weather fluctuations during a period cannot trigger a tipping point, but only if they persist into the next period

tipping and is proportional to $\frac{h(\boldsymbol{X}_t, \boldsymbol{X}_{t+1}(E_t))}{\partial E_t}[V_0 - V_1]$. As long as temperatures are on the rise, increasing emissions increases the future temperature and, thus, increases the likelihood of future tipping. The difference $V_0 - V_1$ multiplies this hazard with the cost of tipping, i.e., the future value loss of the regime shift (evaluated at the given system state).

Lemoine and Traeger (2014) find that either tipping point increases the optimal current carbon tax (or SCC) by approximately 25%, which varies depending on the precise strength of the tipping point and support of the uniform prior. The authors also find that neglecting learning doubles the estimate of the risk premium. The marginal hazard effect, i.e., emissions' increase of the tipping probability by far outweighs the differential welfare impact, i.e., accounting for the policy's marginal welfare impact in the post-tipping regime.[38]

Lemoine and Traeger (2014) only model one tipping point at a time. Lemoine and Traeger (2016a) and Cai et al. (2016), and Cai and Lontzek (2019) model interactions between multiple tipping points. Lemoine and Traeger (2016a) model three tipping points, where one alters the climate system's temperature feedbacks, a second increases the residence time of atmospheric carbon by reflecting warming-induced changes in oceans, soil carbon dynamics, and standing biomass, and a third tipping point increases the convexity of the DICE damage function encapsulating direct economic impacts from warming caused by, e.g., sea level rise resulting from a collapse of one of the major ice sheets, or a weakening of the Northern Atlantic Ocean currents. In Cai et al. (2016) and Cai and Lontzek (2019), all tipping points are represented as a level change of the DICE damage function calibrated to the possible consequences of tipping points in the climate system. A nice novelty of Cai et al. (2016) is that they model a slow delayed onset of the economic impacts of their tipping points in order to be able to better represent crossing a non-convex threshold in the climate system when projecting them straight onto an economic damage function. The solutions of these models are similar to the algorithm suggested above. Lemoine and Traeger (2016a) solve the infinite time-horizon problem recursively in state space for all possible combinations of tipping regimes, starting with the situation where all tipping points have occurred. The social planner eventually solves the present-day policy where she is aware that these irreversible tipping points can occur at any point and in any order, and she understands that her emission and investment decisions affect the tipping hazard. Cai et al. (2016) and Cai and Lontzek (2019) solve the model recursively in finite time, starting from a last period long into the future. Apart from the nature of the tipping points, the main differences between the two approaches are that (i) Lemoine and Traeger (2016a) include learning and (ii) Cai et al. (2016) and Cai and Lontzek (2019) use Epstein-Zin-Weil preferences. Both sets of authors recalibrate DICE's climate system, but Cai et al. (2016) and Cai and Lontzek (2019) retain more state variables and use a supercomputer to solve the model. Interestingly, despite their many differences, Lemoine and Traeger (2016a) and Cai and Lontzek (2019) both find that three interacting tipping points imply a doubling of the optimal present-day carbon tax.[39]

Even more interesting, such a doubling of the optimal carbon tax as a result of tipping points is precisely what Nordhaus (2008) finds as a result of a back-of-the-envelope calculation when calibrating his damage coefficients. Based on a survey about different tipping points, their approximate probabilities, and

[38] We note that this decomposition analyzes the add-on to a given period's carbon tax from the immanent possibility of tipping, which still employs the future value functions. As a result, this decomposition cannot be compared directly to the comprehensive analytic discussion in Section (3), where we recursively eliminate the future value function to explain the full risk premium and account for the interactions of different value function derivatives over time.

[39] Cai and Lontzek (2019) simultaneously include growth uncertainty, which, as discussed in Section 3.6 causes a negative risk premium but increases the SCC based on the calibration effect implied by Epstein-Zin-Weil preferences.

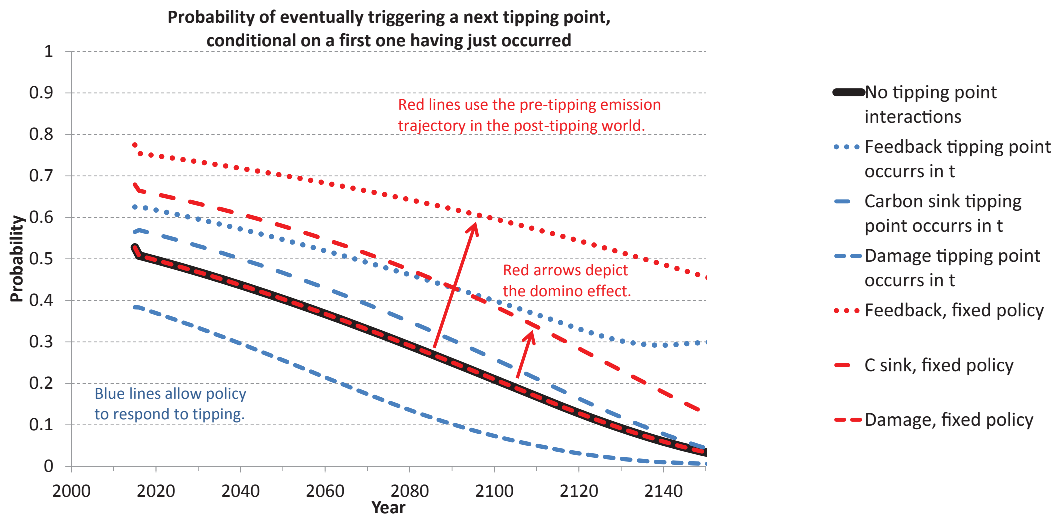

Figure 5: The probability of eventually triggering a next tipping point, conditional on having just triggered a first tipping point (out of three potential ones) in a given year. The black line depicts the case in which tipping points do not interact through system dynamics or policy. The red lines combine the pre-tipping emission trajectory with post-tipping dynamics. The change from the black to the red lines gives each tipping point's (physical) domino effect. The blue lines allow policy and the economy to respond to having tipped. Triggering one of the two climate regime shifts substantially increases the probability of triggering a further tipping point. In contrast, triggering the economic damage tipping point does not increase the probability of triggering another tipping point. The optimal policy response to tipping reduces the domino effect. Replicated from Lemoine and Traeger (2016a).

crude impact guesstimates, Nordhaus approximately doubles the damage coefficients in DICE's damage function. Amazingly, both groups of researchers effectively come up with the same doubling after running through the hoops of complex recursive dynamic programming. In absolute terms, Lemoine and Traeger's (2016a) optimal present-day SCC matches DICE's SCC because the authors removed Nordhaus' tipping point correction to avoid double-counting. Cai and Lontzek (2019) find double the SCC because they retained Nordhaus' original correction in addition to their explicit modeling of tipping points, which is likely closer to today's perspective on the magnitude of climate change damages.

At first sight, we would expect that the inclusion of learning should lead to lower estimates in Lemoine and Traeger (2016a), whereas the use of Epstein-Zin-Weil preferences should lead to a higher estimate in Cai and Lontzek (2019). Figure 5 from Lemoine and Traeger (2016a) provides a possible explanation as to why these two quite different approaches result in a similar correction of the SCC. The figure depicts the probability of eventually triggering the second threshold-crossing after a first threshold has already been crossed. This probability falls over time: if we have not crossed the second tipping point in a later year with a higher temperature, we have expanded the safe temperature range, and, given we are closer to the temperature peak, the probability of eventually tipping has decreased. The red lines show the probability of triggering the second tipping point conditional on holding policy fixed. The thick black line shows the counterfactual scenario where tipping points would not interact. The dotted red line shows the drastic increase in the tipping hazard after the first tipping point has activated the temperature feedbacks. The same emissions now lead to more warming making it more likely to cross the next tipping threshold, a domino effect. The long dashes represent the increased tipping hazard if the first tipping point activates the reduction in CO_2 sequestration from the atmosphere. Also such a reduction in CO_2 sequestration increases the warming and, hence, the likelihood of tipping. Finally, the narrowly dotted line lies right on top of the counterfactual, stating that, given a fixed policy, the damage tipping point does not alter the probability of tipping at all because it does not affect emissions. The corresponding blue lines turn on the optimal policy response. Having tipped implies that the same emissions now lead to either higher warming or higher damages. Either of these implications increases the policy maker's optimal carbon abatement, reducing the tipping hazard as compared to the fixed policy world. If the tipping points alter climate dynamics, the tipping probabilities are still higher also when accounting for the optimal policy response. If the tipping point merely affects economic output, the likelihood of tipping is even smaller after the first tipping point has occurred because higher damages increase optimal abatement and, given the lack of an explicit climate system consequence, an abatement reduction unambiguously reduces the likelihood of tipping. Lemoine and Traeger (2016a) find that the interaction of the temperature feedback tipping point with the damage tipping point contributes the largest "domino" premium to the SCC, followed by the interaction of the carbon sink reduction with the damage tipping point and a triple-interaction premium.

Recognizing that Cai and Lontzek (2019) directly model the economic implications of tipping points, the domino effect in their model is substantially smaller than in Lemoine and Traeger (2016a). In fact, every tipping point occurring in Cai and Lontzek (2019) reduces the likelihood of the next tipping point. Probably, it reduces the likelihood somewhat less than in Figure 5 because the tipping points in Cai and Lontzek (2019) only affect the damage level, whereas the damage tipping point in Lemoine and Traeger (2016a) focuses more on the convexity of damages. The idea of a tipping point increasing damages more at higher temperatures conjectures that, e.g., sea-level rise interacts with the general economic and environmental pressure from global warming. In this case, also the marginal damages from a ton of CO_2 are higher in the post-tipping regime (and more so at higher temperatures). For mere jumps in the damage level, crossing a tipping threshold has severe welfare impacts but triggers less of an ex-post policy response and, thus, less strongly a reduction of the tipping hazard. We note that the absence of learning in Cai and Lontzek (2019) implies that we would not literally be able to plot an analog of Figure 5 in their setting because eventual tipping happens with certainty, and all lines would be unity independent of the policy response. It is probably fair to say that quickly learning whether a tipping point has occurred at a given temperature, as in Lemoine and Traeger (2014), might be the opposite extreme of no learning, and the truth lies somewhere in between (and requires more state variables to model). Related is a scenario by Lemoine and Traeger (2016a), which assumes that the policymaker cannot respond directly to a tipping threshold without delay. They find that the inability to respond directly to a tipping point comes at a substantial welfare loss, which is much larger than a generic policy delay in the absence of tipping points (or when at least coordinating on optimal policy once the first tipping point has occurred).

Cai et al. (2015) introduce environmental amenities as an explicit part of the welfare function and model a tipping point (primarily) affecting these environmental amenities. Given limited substitutability in consumption between environmental amenities and produced goods, such a tipping point can have a particularly strong repercussion on the SCC for an elasticity of substitution below unity, in line with work analyzing direct damages on environmental amenities under limited substitutability in consumption in the DICE model (Sterner and Persson, 2008; Drupp and Hänsel, 2021). Most papers find negligible precautionary savings effects in the face of tipping points (unless they omit economic growth). Van der Ploeg and de Zeeuw (2018) find a more notable precautionary savings effect and discuss that the savings response to a potential tipping point can be positive or negative, comparing stable saddle paths before and after tipping. Their paper also takes the tipping point discussion in IAMs to a continuous time framework that they solve using log-linearization. Also in continuous time, Brausmann and Bretschger's (2024) AK-model sets the stage for modeling tipping points as simple Poisson shocks to economic output, an idea also taken up by Van Den Bremer et al. (2023) in a paper comparing approximate analytic to numeric solutions.

Dietz et al. (2021) integrate all seven tipping point examples mentioned in the introduction into an economic IAM, plus an ice-albedo feedback governing the melting of Arctic sea ice that reduces our planet's reflectivity, further increasing warming. Their calibration of tipping hazards and economic impacts is arguably the most careful we have seen in the literature. Interestingly, their study finds that these eight interactive tipping points increase the SCC by "only" 25%. While their model is not a full-fledged IAM, it accounts for impacts at the country level. Production grows exogenously, but tipping points have immediate and lagged effects on output using a simple lag equation. The fact that policies are exogenous should only increase the impacts of the tipping points on today's SCC. We are not aware of an explanation for why this study differs as much from the earlier publications and try to at least give a possible explanation with respect to Lemoine and Traeger's (2016a) doubling of the SCC. The earlier authors were motivated by a slightly different question. The Copenhaguen and Paris Accord's $2°C$ target was largely based on avoiding dangerous tipping points. Yet, the interaction of tipping points had not previously been analyzed in an integrated climate-economic model. The authors analyze whether, everything else equal, the presence and potential domino-like interaction of tipping points can indeed bring DICE's optimal peak temperature down to $2°C$. For this purpose, the authors consciously picked somewhat extreme tipping point parameterizations and found that, yet, the explicit modeling of tipping points alone will not imply that baseline DICE gives rise to a $2°C$ peak temperature unless we alter other parameters (which very much stands to reason). In fact, the authors found that, all they could achieve was reproducing what Nordhaus had already built into DICE as a damage correction for missing tipping points. The second difference is that Lemoine and Traeger (2016a) explicitly eliminated the earlier tipping point correction of DICE's damage function, which implies a much lower overall SCC in their study. In absolute terms, Dietz et al.'s (2021) Monte Carlo study does indeed find a larger risk premium on the SCC than Lemoine and Traeger (2016a), who somewhat optimistically incorporate optimal policy response and learning. From this perspective, in absolute terms, also Dietz et al. (2021) study somewhat confirms how well our environmental economics Nobel laureate Bill Nordhaus did with his relatively simple back-of-the-envelope adjustment of the damage coefficient in DICE for missing tipping points back in the day.[40]

In a model of damage tipping points, Taconet et al. (2021) find that the main effect of tipping points in IAMs is merely their impact on expected damages. This finding might explain why Nordhaus might have done so well with his back-of-the-envelope correction. From this perspective, it stands to reason to merely incorporate the expected damages of tipping points into a much easier-to-solve model

[40] It is unclear to what degree these adjustments are still part of the most recent versions of DICE. Moreover, neither of these studies uses Epstein-Zin-Weil preferences, which, apart from no learning, can at least partially explain the higher SCC estimates in Cai et al. (2016) and Cai and Lontzek (2019), even in the absence of domino-type interactions.

with smooth damages. Another reason why the simpler "smooth damages" approach is plausible is the extended time span over which climate tipping points gradually unfold their impacts. Even a rapid ice sheet disintegration will take at least several decades, more likely centuries. From this perspective, the main issue of tipping points for economic IAMs is less the speed of the change and more its irreversibility. However, most current climate change and policy scenarios aim to stabilize temperatures at a reasonable degree of warming. In this case, the irreversible nature of climate change is less important for our models because we are not currently foreseeing that we will bring temperatures back to the current or even pre-industrial level at any time in the reasonably near future.

From this perspective, the relevance of explicitly modeling tipping points might be twofold. First, we need such modeling to better estimate the actual impact of tipping points on economic activity and welfare, which requires the explicit integration of tipping points in the climate system with an integrated climate-economic model that can translate the physical tipping processes into output loss. Once we have these estimates and their temperature dependence, it seems reasonable to integrate them into standard IAMs, complementing sectorial bottom-up damage estimates or macroeconomic top-down estimates that will miss out on the impacts of future tipping points. Of course, the tipping points' economic impacts depend directly on the damage functions in the economy, and updates of these non-tipping damage functions will also require us to rerun models like that of Dietz et al. (2021), or models including at least some policy response to update the damage function's tipping component. Second, the irreversible nature of tipping points becomes essential again in models and policy trajectories that foresee a substantial overshooting and that use carbon dioxide removal to substantially lower temperatures again after crossing a peak.

4.4 Ambiguity Aversion & Robust Control

Climate change uncertainty often falls into the realm of deep uncertainty, where probability distributions are neither unique nor well-defined, and objective risk assessments are seldom possible. For example, Figure 3 in Chapter 1 of this handbook illustrates that scientists have derived a range of probability distributions for climate sensitivity using various approaches, including energy balance models, historical data, and large-scale computational simulations. When we modeled growth uncertainty in Jensen and Traeger (2014), we were surprised to find that no reasonable long-term forecasts of economic growth existed for the time horizons required by IAMs. At the time, we turned to the long-run risk model from asset pricing, which had resolved long-standing asset pricing puzzles by positing small but persistent shocks in the consumption process–a clear leap of faith. When Gillingham et al. (2018) faced the same issue in their study, they conducted an expert elicitation (Christensen et al., 2018). Later, for the Rennert et al. (2021) scenarios mentioned in Section 4.1, Müller et al. (2022) produced the most rigorous econometric study to date, estimating a long-run growth process. However, when RFF first proposed the need for a stochastic forecast spanning over a century, one of the authors humorously compared the task to reading tea leaves.

Behaviorally, there is some evidence based on, e.g., Ellsberg's (1961) famous paradox that individuals indeed distinguish between situations with known probabilities and other situations. Ellsberg (1961) famous publication was merely a smart thought experiment based on careful contemplation, suggesting more than just limited rationality. The decision-theoretic literature has axiomatized a variety of models for decision-making under situations with ambiguous probabilities and "hard" or "deep" uncertainty, most of them with a behavioral motivation. From a normative perspective, we might argue that von Neumann and Morgenstern's (1944) axioms still hold the key to how to make rational decisions. Yet, Traeger (2010) shows that even under von Neumann and Morgenstern's (1944) axioms, risk aversion can depend on the level of confidence in probabilistic beliefs once we are willing to distinguish different classes of probabilistic descriptions of the world. The most widely applied ambiguity model in the climate change context falls precisely into this category. The recursive smooth ambiguity framework by Klibanoff et al. (2009) builds a dynamically consistent version of the smooth ambiguity preferences originally axiomatized by Klibanoff et al. (2005). In this framework, ambiguity is represented as a probability distribution. The smooth ambiguity model proposes that decision makers value subjective priors differently from objective risk.

From a different perspective, Nobel laureate Lars Hansen worries about the fact that most of our models are misspecified in one way or another, promoting a robust control approach (Hansen and Sargent, 2001, 2008). As the authors show, this approach can be formally related to maximin expected utility by Gilboa and Schmeidler (1989), which is an extreme limit of smooth ambiguity aversion. However, the perspective of model misspecification differs regarding what can be learned from observation. Robust control and smooth ambiguity aversion have been the most prominent frameworks to assess deep uncertainty in the climate change literature.

Starting with smooth ambiguity aversion, we can interpret such a model in terms of a subjective prior governing the distribution over climate sensitivity and weather shocks that are predicted by well-known probabilities (risk).

$$V(\boldsymbol{X}_t, \boldsymbol{I}_t) = \max_{\boldsymbol{a}_t} \ u(C_t) + \beta \ f^{-1} \left\{ \int f \ \mathbf{E}_t \left[V(\boldsymbol{X}_{t+1}, \boldsymbol{I}_{t+1}) \mid \boldsymbol{X}_t, \boldsymbol{I}_t \right] \mathrm{d}\, \theta_t \right\} \tag{24}$$

Equation (24) modifies the general stochastic dynamic programming equation incorporating an additional aversion to subjective risk. We already wrote the Epstein-Zin-Weil model in equation (17) in a formulation emphasizing the similarity. The difference to the earlier formulation is that equation (24) contains two expectations, and only the subjective prior is evaluated with the additional aversion to risk. Here, θ_t captures this prior uncertainty, and the additional aversion to ambiguity is captured by a concave function f. Temperature shocks or other objective uncertainty are taken care of by the expected value operator $\mathbf{E}_t$ and are not subject to the additional aversion. The decision maker first takes expectations over the future value and then integrates over the subjective prior. In general, the prior evolves over time as the decision maker learns.

Lange and Treich (2008) integrate such a model into a simple stylized two- and three-period framework to sign the effects of ambiguity aversion and learning. Millner et al. (2013) apply the recursive smooth ambiguity aversion model evaluating simulated consumption paths using the DICE model, representing twenty different distributions for climate sensitivity. The decision maker faces ambiguity about which distribution is correct and assigns equal weight to all of them. They examine two policy scenarios from DICE: simulations using business-as-usual emissions and simulations aiming at limiting CO_2 concentrations to twice the preindustrial level. Their results show that ambiguity aversion increases the value of emission reductions. The effect is sizable for a low rate of pure time preference and highly convex damages. Lemoine and Traeger (2016b) use ambiguity aversion to evaluate tipping points, where the decision maker's deep uncertainty governs the distribution of the tipping threshold. In this framework, the decision maker chooses policy optimally and updates a uniform prior governing the location of the tipping point as temperatures rise, learning whether a threshold has been crossed. Numerically, they find that ambiguity aversion leads to a rather moderate increase in emission reductions. Analytically, they discuss how ambiguity aversion affects the *marginal hazard effect* and the *differential welfare impact* discussed in Section 4.3.

Berger et al. (2017) apply a two-period version of the smooth ambiguity preferences developed by Marinacci (2015) in a theoretical exploration of how ambiguity affects optimal emission reductions when there is uncertainty regarding the probability of crossing a 'catastrophic´ tipping point, which would result in a loss of future consumption. In this stylized setting, they identify two conditions under which ambiguity increases optimal abatement. First, the decision maker must exhibit decreasing absolute ambiguity aversion (i.e., ambiguity prudence). Second, the structure of model uncertainty must be such that the disagreement between models does not increase as abatement increases.[41] The authors numerically analyze the impact of smooth ambiguity aversion using their own adaptation of the DICE model. They introduce a tipping point that could occur in the year 2100, at which time the decision maker learns exogenously whether the 'catastrophe´ - a 20 percent reduction in GDP - has occurred. To evaluate welfare, they use preferences that disentangle risk, ambiguity, and intertemporal substitution. Their results show that, in this setting, ambiguity aversion increases abatement by approximately one-fourth to one-half of the effect of risk aversion.

We now turn to robust control, the second prominent approach addressing deep uncertainty in climate economics. It is worth noting that the majority of robust control problems are formulated in continuous time, and often multiple alternative mathematical formulations exist. Whether one uses discrete or continuous time is generally a matter of convenience; for some specific problems, the analysis is more straightforward in one format than the other. As mentioned above, this approach is often motivated by model misspecification and, therefore, also called model uncertainty. In this framework, the decision maker has a preferred or 'best-guess' model but also considers a set of alternative models that might be correct. This set consists of models that are within a certain 'distance' from the preferred one.[42] In an ambiguity context, the size of this set can relate to both the level of ambiguity (i.e., how much the model might be misspecified) and the degree of aversion to this uncertainty. When choosing a policy, the decision maker accounts for the possibility that her preferred model could be wrong.

Let p_t denote the probability distribution governing the decision maker's best-guess model, and let $D_t(p_{t+1}|\, q_{t+1})$ measure the distance between this best-guess model and an alternative, possibly correct model q_{t+1}. Here, the "model" is captured by the probability distributions governing our equations of motion.

A helpful way to think about this framework is by considering a fictitious adversary, often referred to as 'nature.' Nature aims to minimize the decision maker's welfare. We can formulate a game between the decision maker maximizing welfare by choosing a_t, and nature choosing p_t to minimize the payoff

$$V(\boldsymbol{X}_t, \boldsymbol{I}_t) = \min_{q_t} \left(\max_{\boldsymbol{a}_t} \quad u(C_t) + \beta \left\{ \int \left[\psi D_{t+1}(p_{t+1}|\, q_{t+1}) + \; V(\boldsymbol{X}_{t+1}, \boldsymbol{I}_{t+1}) \right] dq_{t+1} \right\} \right). \tag{25}$$

Equation (25) is adapted from Rudik (2020). Being precautions, the decision maker evaluates the future for all possible models q_{t+1}. She then assumes somewhat pessimistically that nature draws the worst of

[41] Alternatively, if the decision maker exhibits constant absolute ambiguity aversion, the uncertainty must decrease with greater abatement.

[42] Formally, this distance is measured using an entropy metric, often referred to as an entropy ball.

these possible futures with one twist: the term $\psi D_{t+1}(p_{t+1}|\, q_{t+1})$ increases the value under the integral for models that are far from the best-guess. As a result, nature will be less prone to pick a model with moderately lower payoffs if it is far away from the best-guess model. The nonnegative parameter ψ is often called the penalty parameter.

Hennlock (2009) and Athanassoglou and Xepapadeas (2012) were among the first to apply robust control preferences in the context of climate change analysis. Hennlock (2009) introduces uncertainty around equilibrium climate sensitivity and derives an analytic solution to his model by making several simplifying assumptions, such as linear damages and an additively separable objective function. His findings suggest that robust control preferences generally lead to increased abatement efforts. Athanassoglou and Xepapadeas (2012), on the other hand, utilize a linear-quadratic utility framework and simple carbon stock dynamics to introduce uncertainty. These assumptions allow them to obtain an analytical solution to a partial equilibrium problem. While their theoretical results show that increased uncertainty does not necessarily result in greater emission reductions, their numerical calibrations consistently indicate that uncertainty leads to higher abatement.

Anderson et al. (2014) use a simplified version of Golosov et al.'s (2014) model with only one energy source and improve the climate model by using the TCRE model discussed in Section 2.1. In addition to providing analytic solutions, they also solve a more complex version of their model numerically. Their model incorporates three simultaneous uncertainties: temperature, damages, and total factor productivity. One of their key insights is that introducing robust control preferences can either increase or decrease optimal emissions, depending on whether the consumption smoothing parameter is greater or less than one. Li et al. (2016) focus on damage uncertainty, also using a model based on Golosov et al. (2014). They find that increasing the decision maker's concern about misspecification leads to a higher optimal carbon tax, a concern that increases with carbon emissions.

Barnett et al. (2020) develop a framework for climate economics that addresses model misspecification and model uncertainty separately—capturing ambiguity both *within* models and *between* models. Using decision-theoretic and asset pricing tools, their model incorporates the preference specification by Hansen and Miao (2018), which separates aversion to model misspecification and model uncertainty into distinct parameters. Much like standard robust control preferences, there is a cost to deviating from the preferred model specification, both within and between models. The decision maker is uncertain about climate sensitivity and the damage function, with damages either affecting output or growth. They find that uncertainty governing climate sensitivity and damages is multiplicative. The authors show how uncertainty preferences can be expressed as adjustments to probabilities. Their numerical illustrations suggest that model uncertainty has a modest impact on the optimal social cost of carbon when damages reduce GDP but a much larger impact when damages reduce growth. Barnett et al. (2022) extend the framework and analysis in Barnett et al. (2020), introducing uncertainty in the carbon cycle and decomposing the contributions of the three uncertainties to the optimal carbon tax.

Rudik (2020) introduces uncertainty in the damage parameters into a stochastic dynamic programming version of the DICE model and evaluates it using robust control preferences. By analytically decomposing the optimal carbon tax using a Taylor expansion, Rudik (2020) identifies a channel he refers to as "misspecification insurance," which affects the tax. He finds that this channel increases the optimal carbon tax in the absence of learning but decreases it when learning is introduced. His numerical results suggest that robust control preferences do not significantly alter the optimal tax.

4.5 Approximate Dynamic Programming and Machine Learning

This final section slightly differs from the preceding sections with what we consider a motivational methodological outlook into the future. Given the recent success of and interest in machine learning approaches, we want to briefly motivate how these techniques can be made useful for dynamic programming and the challenge of solving more complex IAMs numerically. In the climate change context, this approach is pioneered by Friedl et al. (2023). In our discussion of classical dynamic programming, we focused on solving the Bellman equation by iterating on the value function, where the right-hand side of the Bellman equation (22) is evaluated and used to update the left-hand side. This iterative process often relies on contraction mappings and function approximation techniques like Chebyshev polynomials. However, this method can become computationally expensive, especially in high-dimensional state spaces typical in climate change economics. Alternatively, we can directly iterate on the Euler equation. The "machine learning" approach picks up ideas from approximate dynamic programming using such a forward simulation (Bertsekas et al., 2011), neural networks for function approximation, and stochastic gradient descent for parametric updates of the approximation.

The forward simulation approach involves generating simulated paths of the state variables over time rather than iterating directly on the Bellman equation. These simulated paths allow us to explore the dynamics of the system under uncertainty, often using Monte Carlo methods. Monte Carlo simulation involves generating multiple possible future trajectories of the state variables, given initial conditions and stochastic elements like random shocks to the economy or the climate system. The simulation uses the current guess of the control rule for, e.g., consumption and emission policy. The control rules can be approximated by the same neural network that is used to estimate the value function. This is a significant departure from the approach we discussed previously, where the control variables are optimized explicitly at each step or iteration of the Bellman equation.[43]

We can simulate forward using a policy function that the neural network approximates alongside the value function. By jointly approximating both, the neural network allows us to simulate the system dynamics without re-solving the optimization problem at each step, leading to significant computational savings. However, to accurately capture the uncertainties in the problem, we must simulate an event tree that branches out over time with sufficient density to reflect the stochastic nature effectively, which can require breaking up the forward simulation into different batches (see also Footnote 43). To further enhance the efficiency of forward simulation, directly applying the Euler equations governing the optimal controls can be advantageous. The Euler equations characterize the optimal trade-off between current and future utility, offering a direct method for updating the policy function, which can lead to faster convergence and more accurate approximations.

For each simulated trajectory, we compute the immediate utility and an estimate of the future value function based on the current policy or control strategy. The difference between the current estimate of the value function and the value function estimated from the forward simulation provides the basis for updating the model. This approach differs from the traditional Bellman iteration by not requiring the direct computation of the value function at every possible state, thus reducing the computational burden.

[43] Howard's method, mentioned above in the context of classical dynamic programming, similarly avoids re-optimization at every iteration. However, because the classical method solves the problem at every grid point–potentially using a set of Gauss-Legendre points for stochastic realizations–it does not require Monte Carlo forward simulation. Even if Monte Carlo methods are used instead of Gauss-Legendre grid points, the classical approach only requires simulation for one period at a time, thereby avoiding the extensive branching out of the decision tree that occurs in the multi-period forward simulation employed in approximate dynamic programming.

Stochastic Gradient Descent is the optimization algorithm used to update the parameters of the neural network. Its goal is to minimize the error between the current estimate of the value function and the observed outcomes from the forward simulation, which in this economic context is the difference between current and future utility. The temporal difference error, which drives the updates, is computed as

$$\delta_t = u(C_t) + \beta V_\theta(\boldsymbol{X}_{t+1}) - V_\theta(\boldsymbol{X}_t),$$

where $\boldsymbol{X}_{t+1}$ is the next state obtained from the forward simulation. The parameters θ describe the parameters defining the neural network (the value function approximation) and are updated using the gradient of this error

$$\theta_{t+1} = \theta_t - \alpha_t \nabla_\theta \delta_t,$$

where α_t is the learning rate. The learning rate controls the size of the steps the algorithm takes when updating the neural network's parameters. If the learning rate is set too high, the algorithm may take steps that are too large, causing it to overshoot the optimal solution. This can lead to instability or divergence, where the algorithm fails to converge to a solution. On the other hand, if the learning rate is set too low, the algorithm will take very small steps, resulting in slow convergence and potentially requiring an excessive number of iterations to reach a good solution.[44]

Neural networks in these models are typically composed of multiple layers of interconnected neurons, known as deep neural networks. Each neuron in the network can be thought of as a simple function that takes inputs and produces an output, contributing to the network's overall "decision-making process".[45] The network consists of an input layer, several hidden layers, and an output layer. The input layer takes in the state variables $\boldsymbol{X}_t$, while the hidden layers transform these inputs through weighted sums and nonlinear activation functions. These transformations enable the network to capture complex, nonlinear relationships between the state variables and the value function. The output layer provides the estimated value function $V_\theta(s_t)$ or the policy decision. The depth of the network (i.e., the number of hidden layers) allows it to model intricate patterns and dependencies that would be difficult to capture with simpler models.

As we discussed above, a crucial challenge in dynamic programming is the curse of dimensionality, where the computational complexity can grow exponentially (or at a high polynomial rate) with the number of state variables. While neural networks do not universally solve this problem, they offer practical advantages in many cases. Specifically, deep neural networks can learn to represent high-dimensional data in a more compact form through hierarchical or multiscale representations, which can help mitigate the effects of dimensionality. Additionally, under certain conditions, such as in the approximation of smooth functions with bounded derivatives (as shown in Barron's theorem), the number of parameters required by a neural network can grow only polynomially with dimensionality rather than exponentially. In general, neural networks seem capable of approximating complex functions with far fewer parameters than traditional grid-based methods.

[44] To check whether the learning rate is well-chosen, you can monitor the convergence behavior of the algorithm. If the loss function (or error measure) decreases steadily and smoothly, the learning rate is likely appropriate. If the loss function fluctuates wildly or increases, the learning rate may be too high. Conversely, if the loss function decreases very slowly or appears to plateau early, the learning rate may be too low. In practice, one can experiment with different learning rates or use adaptive methods like Adam, which adjust the learning rate dynamically.

[45] More specifically, a neuron in a neural network takes one or more inputs, applies a weighted sum to these inputs, adds a bias term, and then passes the result through a nonlinear activation function.

Speed and convergence to local optima might be the most important characteristics setting the method apart from classical dynamic programming. One of the major strengths of the approach is its computational efficiency. Modern implementations of neural networks leverage highly optimized libraries like TensorFlow or PyTorch, which are designed to perform the necessary tensor algebra operations at high speed. These operations can be parallelized and executed on GPUs, significantly reducing computation time compared to CPU-based implementations. This makes the approach not only scalable to large problems but also fast, allowing for the analysis of more complex climate-economic models in far less time than required by traditional methods.

While the approach offers significant advantages, it is important to acknowledge that there is no underlying theory guaranteeing global convergence. Stochastic gradient descent is a local optimization method, meaning it is possible and also common to converge to a local optimum rather than the global optimum. However, with careful tuning of the learning rate, network architecture, and training process, it is often possible to achieve satisfactory solutions that are close to the global optimum. Experience has shown that the approach often produces high-quality solutions. The flexibility and power of neural networks, combined with their efficiency in connection to simulation-based stochastic gradient descent, make this approach one of the most promising methods for exploring the solution space of high-dimensional dynamic programming problems, even though it does not guarantee the absolute best solution. However, as many of us have seen with large language models that use related methods, there is a risk of what is commonly referred to as "hallucinations." While the hope is that well-defined economic problems with sufficiently smooth objective functions will yield more reliable results in economics, it remains crucial to carefully validate the outcomes when using these methods.

5 Concluding Summary

This chapter has underscored the profound influence of uncertainty on climate policy, particularly in shaping the cost of CO_2 emissions and guiding optimal decision-making. Several key insights have emerged from both numeric and analytic modeling. Unfortunately, one such insight is that there is no simple, uniform rule on how to respond to uncertainty in the climate context. Growth uncertainty affects the optimal carbon price differently than climate uncertainty, and studies often disagree on even the direction of the overall risk premium. For this reason, we have devoted considerable attention to discussing the individual drivers of policy risk premia, and then discussed the models that evaluate their relative importance. We explained the core assumptions of various modeling approaches, such as Monte Carlo simulations versus models of decision-making under uncertainty and learning.

We have shown that the impact of uncertainty on optimal policy depends on how it interacts with time preference, risk preferences, and the various convexities and correlations within climate-economic models. To summarize a few key insights: time preference and uncertainty strongly interact, especially with respect to the higher moments of probability distributions. A related important result is that the risk premium increases convexly in relation to components of the deterministic SCC, making uncertainty more significant when more is at stake.

As is well known, a higher elasticity of intertemporal substitution increases the deterministic SCC (by lowering growth discounting). By contrast, it reduces the risk premium, which can even turn negative. Disentangling the elasticity of intertemporal substitution from risk aversion helps clarify the distinct roles these preference characteristics play in shaping climate policy. This distinction is also critical for more accurate calibration of climate-economic models to both risk-free discount rates and risk premia.

It is crucial to note that it is not risk preferences alone, but the interaction between various model convexities, that plays a decisive role in determining the policy response to risk. Classical prudence arguments from the precautionary savings literature appear to play a more minor role in the climate context. It is moreover important to capture that climate risk is largely endogenous–driven by human activity. Additionally, we expect to learn more about the true climate sensitivity over time. While the precise impact of learning varies across models, we feel comfortable ruling out that wait-and-see is a reasonable response to anticipated learning.

The influence of model uncertainty and ambiguity aversion also varies depending on the model and approach chosen. A potential conjecture is that the results may not differ as much from those under an 'objective risk' framework as one might expect. The interaction of different uncertainties often amplifies the need for stronger policy responses. This is true for both smooth, reversible uncertainties and potential tipping points. While the direct impact of irreversible damages can be captured through an increase in expected damages, translating the potential interactions between tipping points into economic damages requires more complex climate-economic models.

This chapter has explored the strengths and limitations of various modeling approaches, each offering valuable insights into climate-economic uncertainties and their implications for current policy. Analytic models help identify the drivers of the uncertainty premium, while numeric models assess their relative importance. We hope that new machine learning-based solution methods will enable the analysis of even more complex stochastic interactions, guided by the analytic insights we have discussed. One particularly promising area for future research is the interaction of uncertainty across heterogeneous regions–an issue we hope to see addressed in a future handbook chapter.

Declaration of AI and AI-assisted technologies in the writing process: During the preparation of this work, the authors used ChatGPT and Grammarly to improve the writing style. After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

Allen, M. R., D. J. Frame, C. Huntingford, C. D. Jones, J. A. Lowe, M. Meinshausen, and N. Meinshausen (2009). Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature 458*(7242), 1163–1166.

Anderson, E., W. Brock, L. P. Hansen, and A. H. Sanstad (2014). Robust analytical and computational explorations of coupled economic-climate models with carbon-climate response. *RDCEP Working Paper 13-05*.

Anthoff, D. and R. S. Tol (2013). The uncertainty about the social cost of carbon: A decomposition analysis using fund. *Climatic change 117*, 515–530.

Armstrong McKay, D. I., A. Staal, J. F. Abrams, R. Winkelmann, B. Sakschewski, S. Loriani, I. Fetzer,

S. E. Cornell, J. Rockström, and T. M. Lenton (2022). Exceeding 1.5 c global warming could trigger multiple climate tipping points. *Science 377*(6611), eabn7950.

Athanassoglou, S. and A. Xepapadeas (2012). Pollution control with uncertain stock dynamics: when, and how, to be precautious. *Journal of Environmental Economics and Management 63*(3), 304–320.

Bansal, R., D. Kiku, and M. Ochoa (2019). Climate change risk. *Federal Reserve Bank of San Francisco Working Paper*.

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron (2014). Volatility, the macroeconomy, and asset prices. *The Journal of Finance 69*(6), 2471–2511.

Bansal, R., D. Kiku, and A. Yaron (2010). Long-run risks, the macro-economy and asset prices. *American Economic Review: Papers & Proceedings 100*, 542–546.

Bansal, R., D. Kiku, and A. Yaron (2012). An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review 1*, 183–221.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance 59*(4), 1481–509.

Barnett, M., W. Brock, and L. P. Hansen (2020). Pricing uncertainty induced by climate change. *The Review of Financial Studies 33*(3), 1024–1066.

Barnett, M., W. Brock, and L. P. Hansen (2022). Climate change uncertainty spillover in the macroeconomy. *NBER Macroeconomics Annual 36*(1), 253–320.

Barrage, L. and W. Nordhaus (2024). Policies, projections, and the social cost of carbon: Results from the dice-2023 model. *Proceedings of the National Academy of Sciences 121*(13), e2312030121.

Bauer, M. D. and G. D. Rudebusch (2023). The rising cost of climate change: evidence from the bond market. *Review of Economics and Statistics 105*(5), 1255–1270.

Bellman, R. (1966). Dynamic programming. *science 153*(3731), 34–37.

Berger, L., J. Emmerling, and M. Tavoni (2017). Managing catastrophic climate risks under model uncertainty aversion. *Management Science 63*(3), 749–765.

Berger, L. and M. Marinacci (2020). Model uncertainty in climate change economics: A review and proposed framework for future research. *Environmental and Resource Economics*, 1–27.

Bertsekas, D. P. et al. (2011). Dynamic programming and optimal control 3rd edition, volume ii. *Belmont, MA: Athena Scientific 1*.

Bosetti, V., C. Carraro, M. Galeotti, E. Massetti, and M. Tavoni (2006). A world induced technical change hybrid model. *The Energy Journal 27*(2_suppl), 13–37.

Brausmann, A. and L. Bretschger (2024). Escaping damocles' sword: Endogenous climate shocks in a growing economy. *Environmental and Resource Economics*, 1–48.

Brock, W. and D. Starrett (2003, December). Managing systems with non-convex positive feedback. *Environmental and Resource Economics 26*(4), 575–602.

Burke, M., S. M. Hsiang, and E. Miguel (2015). Global non-linear effect of temperature on economic production. *Nature 527*(7577), 235–239.

Cai, Y. (2021). The role of uncertainty in controlling climate change. In *Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press. Published online: 23 February 2021.

Cai, Y., K. L. Judd, T. M. Lenton, T. S. Lontzek, and D. Narita (2015). Environmental tipping points significantly affect the cost- benefit assessment of climate policies. *Proceedings of the National Academy of Sciences 112*(15), 4606–4611.

Cai, Y., T. M. Lenton, and T. S. Lontzek (2016). Risk of multiple interacting tipping points should encourage rapid co2 emission reduction. *Nature Climate Change 6*(5), 520–525.

Cai, Y. and T. S. Lontzek (2019). The social cost of carbon with economic and climate risks. *Journal of Political Economy 127*(6), 2684–2734.

Calvin, K., P. Patel, L. Clarke, G. Asrar, B. Bond-Lamberty, M. Hejazi, P. Kyle, M. Wise, A. Di Vittorio, and J. Edmonds (2011). *GCAM Documentation*. Joint Global Change Research Institute, College Park, MD.

Carleton, T. and M. Greenstone (2022). A guide to updating the us government's social cost of carbon. *Review of Environmental Economics and Policy 16*(2), 196–218.

Carleton, T., A. Jina, M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, R. E. Kopp, K. E. McCusker, I. Nath, et al. (2022). Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *The Quarterly Journal of Economics 137*(4), 2037–2105.

Chen, X., J. Favilukis, and S. C. Ludvigson (2013). An estimation of economic models with recursive preferences. *Quantitative Economics 4*, 39–83.

Christensen, P., K. Gillingham, and W. Nordhaus (2018). Uncertainty in forecasts of long-run economic growth. *Proceedings of the National Academy of Sciences 115*(21), 5409–5414.

Clarke, H. R. and W. J. Reed (1994, September). Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse. *Journal of Economic Dynamics and Control 18*(5), 991–1010.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2016). Parameter Learning in General Equilibrium: The Asset Pricing Implications. *American Economic Review 106*(3), 664–698.

Costello, C. J., M. G. Neubert, S. A. Polasky, and A. R. Solow (2010). Bounded uncertainty and climate change economics. *Proceedings of the National Academy of Sciences 107*(18), 8108–8110.

Crost, B. and C. P. Traeger (2013). Optimal climate policy: Uncertainty versus Monte-Carlo. *Economic Letters 120*(3), 552–558.

Crost, B. and C. P. Traeger (2014). Optimal co2 mitigation under damage risk valuation. *Nature Climate Change 4*, 631–636.

Daniel, K. D., R. B. Litterman, and G. Wagner (2019). Declining co2 price paths. *Proceedings of the National Academy of Sciences 116*(42), 20886–20891.

Dekel, E., B. L. Lipman, and A. Rustichini (1998). Recent developments in modeling unforeseen contingencies. *European Economic Review 42*(3-5), 523–542.

Denning, A. S. (2022). Where has all the carbon gone? *Annual Review of Earth and Planetary Sciences 50*(1), 55–78.

Dietz, S., C. Gollier, and L. Kessler (2018). The climate beta. *Journal of Environmental Economics and Management 87*, 258–274.

Dietz, S., J. Rising, T. Stoerk, and G. Wagner (2021). Economic impacts of tipping points in the climate system. *Proceedings of the National Academy of Sciences 118*(34), e2103081118.

Dietz, S., F. van der Ploeg, A. Rezai, and F. Venmans (2021). Are economists getting climate dynamics right and does it matter? *Journal of the Association of Environmental and Resource Economists 8*(5), 895–921.

Dietz, S. and F. Venmans (2019). Cumulative carbon emissions and economic policy: in search of general principles. *Journal of Environmental Economics and Management 96*, 108–129.

Drupp, M. A. and M. C. Hänsel (2021). Relative prices and climate policy: how the scarcity of nonmarket goods drives policy evaluation. *American Economic Journal: Economic Policy 13*(1), 168–201.

Ellsberg, D. (1961). Risk, ambiguity and the savage axioms. *Quarterly Journal of Economics 75*, 643–69.

EPA (2023). Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances.

Epstein, L. G., E. Farhi, and T. Strzalecki (2014, 09). How much would you pay to resolve long-run risk? *The American Economic Review 104*, 2680–2697.

Epstein, L. G. and S. E. Zin (1989, July). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*(4), 937–69.

Epstein, L. G. and S. E. Zin (1991, April). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *Journal of Political Economy 99*(2), 263–86.

Fagin, R. and J. Y. Halpern (1987). Belief, awareness, and limited reasoning. *Artificial intelligence 34*(1), 39–76.

Folini, D., A. Friedl, F. Kübler, and S. Scheidegger (2024). The climate in climate economics. *Review of Economic Studies*, rdae011.

Friedl, A., F. Kübler, S. Scheidegger, and T. Usui (2023). Deep uncertainty quantification: with an application to integrated assessment models. Technical report, Working Paper University of Lausanne.

Geoffroy, O., D. Saint-Martin, D. J. Olivié, A. Voldoire, G. Bellon, and S. Tytéca (2013). Transient climate response in a two-layer energy-balance model. part i: Analytical solution and parameter calibration using cmip5 aogcm experiments. *Journal of climate 26*(6), 1841–1857.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*(2), 141–53.

Gillingham, K., W. Nordhaus, D. Anthoff, G. Blanford, V. Bosetti, P. Christensen, H. McJeon, and J. Reilly (2018). Modeling uncertainty in integrated assessment of climate change: A multimodel comparison. *Journal of the Association of Environmental and Resource Economists 5*(4), 791–826.

Gollier, C. (2004). Maximizing the expected net future value as an alternative strategy to gamma discounting. *Finance Research Letters 1*(2), 85–89.

Gollier, C. (2013). *Pricing the planet's future: the economics of discounting in an uncertain world.* Princeton University Press.

Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica 82*(1), 41–88.

Golub, A., D. Narita, and M. G. Schmidt (2014). Uncertainty in integrated assessment models of climate change: Alternative analytical approaches. *Environmental Modeling & Assessment 19*(2), 99–109.

Ha-Duong, M. and N. Treich (2004). Risk aversion, intergenerational equity and climate change. *Environmental and Resource Economics 28*(2), 195–207.

Hänsel, M. C., M. A. Drupp, D. J. Johansson, F. Nesje, C. Azar, M. C. Freeman, B. Groom, and T. Sterner (2020). Climate economics support for the un climate targets. *Nature Climate Change 10*(8), 781–789.

Hansen, L. P. and J. Miao (2018). Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proceedings of the National Academy of Sciences 115*(37), 9163–9168.

Hansen, L. P. and T. J. Sargent (2001). Robust control and model uncertainty. *American Economic Review 91*(2), 60–66.

Hansen, L. P. and T. J. Sargent (2008). *Robustness.* Princeton university press.

Heal, G. (1984). Interactions between economy and climate: A framework for policy design under uncertainty. In V. K. Smith and A. D. White (Eds.), *Advances in Applied Microeconomics*, Volume 3, pp. 151–168. Greenwich, CT: JAI Press.

Heal, G. (2017). The economics of the climate. *Journal of Economic Literature 55*(3), 1046–1063.

Heifetz, A., M. Meier, and B. C. Schipper (2013). Unawareness, beliefs, and speculative trade. *Games and Economic Behavior 77*(1), 100–121.

Hennlock, M. (2009). Robust control in global warming management: An analytical dynamic integrated assessment. *RFF Workding Paper* (DP 09-19).

Heutel, G., J. Moreno-Cruz, and S. Shayegh (2018). Solar geoengineering, uncertainty, and the price of carbon. *Journal of Environmental Economics and Management 87*, 24–41.

Hoel, M. and L. Karp (2001). Taxes and quotas for a stock pollutant with multiplicative uncertainty. *Journal of Public Economics 82*, 91–114.

Hoel, M. and L. Karp (2002). Taxes versus quotas for a stock pollutant. *Resource and Energy Economics 24*, 367–384.

Hope, C. (2006). The marginal impact of CO2 from PAGE2002: An integrated assessment model incorporating the IPCC's five reasons for concern. *The Integrated Assessment Journal 6*(1), 19–56.

Howard, P. H. and T. Sterner (2017). Few and not so far between: A meta-analysis of climate damage estimates. *Environmental and Resource Economics 68*, 197–225.

Hsiang, S., R. Kopp, A. Jina, J. Rising, M. Delgado, S. Mohan, D. J. Rasmussen, R. Muir-Wood, P. Wilson, M. Oppenheimer, et al. (2017). Estimating economic damage from climate change in the united states. *Science 356*(6345), 1362–1369.

Hwang, I. C., F. Reynès, and R. S. Tol (2017). The effect of learning on climate policy under fat-tailed risk. *Resource and Energy Economics 48*, 1–18.

Interagency Working Group on Social Cost Greenhouse Gases, U. S. G. (2021). Technical support document: Social cost of carbon, methane, and nitrous oxide interim estimates under executive order 13990. *Department of Energy*.

Interagency Working Group on Social Cost of Carbon, U. S. G. (2010). Technical support document: Social cost of carbon for regulatory impact analysis under executive order 12866. *Department of Energy*.

Interagency Working Group on Social Cost of Carbon, U. S. G. (2013). Technical support document: Technical update of the social cost of carbon for regulatory impact analysis under executive order 12866. *Department of Energy*.

Interagency Working Group on Social Cost of Carbon, U. S. G. (2016). Technical support document: Technical update of the social cost of carbon for regulatory impact analysis under executive order 12866. *Department of Energy*.

IPCC (2013). *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.

IPCC (2021). *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.

Jensen, S. and C. Traeger (2013). Optimally climate sensitive policy: A comprehensive evaluation of uncertainty & learning. *Department of Agricultural and Resource Economics, UC Berkeley*.

Jensen, S. and C. P. Traeger (2014). Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *European Economic Review 69*(C), 104–125.

Jensen, S. and C. P. Traeger (2021). Pricing climate risk.

Joos, F., R. Roth, and A. J. Weaver (2013). Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis. *Atmos. Chem. Phys. 13*, 2793–2825.

Judd, K. L. (1996). Approximation, perturbation, and projection methods in economic analysis. *Handbook of computational economics 1*, 509–585.

Karp, L. and C. Traeger (2024). Smart cap. *Journal of the European Economic Association*, jvae030.

Karp, L. and J. Zhang (2006). Regulation with anticipated learning about environmental damages. *Journal of Environmental Economics and Management 51*, 259–279.

Keller, K., B. M. Bolker, and D. F. Bradford (2004). Uncertain climate thresholds and optimal economic growth. *Journal of Environmental Economics and Management 48*, 723–741.

Kelly, D. L., G. Heutel, J. B. Moreno-Cruz, and S. Shayegh (2024, November). Solar geoengineering, learning, and experimentation. *Journal of the Association of Environmental and Resource Economists 11*(6).

Kelly, D. L. and C. D. Kolstad (1999). Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control 23*, 491–518.

Kelly, D. L. and C. D. Kolstad (2001). Solving infinite horizon growth models with an environmental sector. *Computational Economics 18*, 217–231.

Kelly, D. L. and Z. Tan (2015). Learning and climate feedbacks: Optimal climate insurance and fat tails. *Journal of Environmental Economics and Management 72*, 98–122.

Kimball, M. S. (1990). Precautionary saving in the small and in the large. *Econometrica 58*(1), 53–73.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica 73*(6), 1849–1892.

Klibanoff, P., M. Marinacci, and S. Mukerji (2009). Recursive smooth ambiguity preferences. *Journal of Economic Theory 144*, 930–976.

Kolmogoroff, A. (1933). Grundbegriffe der wahrscheinlichkeitsrechnung.

Lange, A. and N. Treich (2008). Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions. *Climatic Change 89*, 7–21.

Leach, A. J. (2007). The climate change learning curve. *Journal of Economic Dynamics and Control 31*, 1728–1752.

Lemoine, D. (2021). The climate risk premium: how uncertainty affects the social cost of carbon. *Journal of the Association of Environmental and Resource Economists 8*(1), 27–57.

Lemoine, D. and C. Traeger (2016a). Economics of tipping the climate dominoes. *Nature Climate Change 6*, 514–519.

Lemoine, D. and C. P. Traeger (2014). Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy 6*(1), 1–31.

Lemoine, D. and C. P. Traeger (2016b). Ambiguous tipping points. *Journal of Economic Behavior & Organization 132*, 5–18.

Li, X., B. Narajabad, and T. Temzelides (2016). Robust dynamic energy use and climate change. *Quantitative Economics 7*(3), 821–857.

Manne, A., R. Mendelsohn, and R. Richels (1995). Merge: A model for evaluating regional and global effects of ghg reduction policies. *Energy policy 23*(1), 17–34.

Marinacci, M. (2015). Model uncertainty. *Journal of the European Economic Association 13*(6), 1022–1100.

Matthews, H. D., N. P. Gillett, P. A. Stott, and K. Zickfeld (2009). The proportionality of global warming to cumulative carbon emissions. *Nature 459*(7248), 829–832.

Meier, F. D. and C. P. Traeger (2023). Uncertain remedies to fight uncertain consequences: The case of solar geoengineering.

Millner, A. (2013). On welfare frameworks and catastrophic climate risks. *Journal of Environmental Economics and Management 65*(2), 310–325.

Millner, A., S. Dietz, and G. Heal (2013). Scientific ambiguity and climate policy. *Environmental and Resource Economics 55*, 21–46.

Moore, F. C., M. A. Drupp, J. Rising, S. Dietz, I. Rudik, and G. Wagner (2024). Synthesis of evidence yields high social cost of carbon due to structural model variation and uncertainties. Technical report, CESifo.

Müller, U. K., J. H. Stock, and M. W. Watson (2022). An econometric model of international growth dynamics for long-horizon forecasting. *Review of Economics and Statistics 104*(5), 857–876.

Mäler, K., A. Xepapadeas, and A. de Zeeuw (2003, December). The economics of shallow lakes. *Environmental and Resource Economics 26*(4), 603–624.

Nakamura, E., D. Sergeyev, and J. Steinsson (2017, January). Growth-rate and uncertainty shocks in consumption: Cross-country evidence. *American Economic Journal: Macroeconomics 9*(1), 1–39.

Nakamura, E., J. Steinsson, R. Barro, and J. Ursua (2013). Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics 5*(3), 35–74.

Newell, R. G., W. A. Pizer, and B. C. Prest (2022). A discounting rule for the social cost of carbon. *Journal of the Association of Environmental and Resource Economists 9*(5), 1017–1046.

Nordhaus, W. (2008). *A question of balance: Weighing the options on global warming policies.* Yale University Press.

Nordhaus, W. and P. Sztorc (2013). Dice 2013r: Introduction and user's manual.

Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences 114*(7), 1518–1523.

Pindyck, R. S. (2007). Uncertainty in environmental economics. *Review of environmental economics and policy*.

Pizer, W. A. (1999). The optimal choice of climate change policy in the presence of uncertainty. *Resource and Energy Economics 21*(3-4), 255–287.

Rennert, K., F. Errickson, B. C. Prest, L. Rennels, R. G. Newell, W. Pizer, C. Kingdon, J. Wingenroth, R. Cooke, B. Parthum, et al. (2022). Comprehensive evidence implies a higher social cost of co2. *Nature 610*(7933), 687–692.

Rennert, K., F. Errickson, B. C. Prest, L. Rennels, R. G. Newell, W. Pizer, C. Kingdon, J. Wingenroth, R. Cooke, B. Parthum, D. Smith, K. Cromar, D. Diaz, F. C. Moore, U. K. Müller, R. J. Plevin, A. E. Raftery, H. Ševčíková, H. Sheets, J. H. Stock, T. Tan, M. Watson, T. E. Wong, and D. Anthoff (2022). Comprehensive evidence implies a higher social cost of co2. *Nature 610*(7933), 687–692.

Rennert, K., B. C. Prest, W. A. Pizer, R. G. Newell, D. Anthoff, C. Kingdon, L. Rennels, R. Cooke, A. E. Raftery, H. Ševčíková, and F. Errickson (2021, October). The social cost of carbon: Advances in long-term probabilistic projections of population, gdp, emissions, and discount rates. Working Paper 21-28, Resources for the Future.

Rode, A., T. Carleton, M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, A. Jina, R. E. Kopp, K. E. McCusker, I. Nath, J. Rising, and J. Yuan (2021). Estimating a social cost of carbon for global energy consumption. *Nature 598*(7880), 308–314.

Rudik, I. (2020). Optimal climate policy when damages are unknown. *American Economic Journal: Economic Policy 12*(2), 340–73.

Rudik, I. and D. Lemoine (2017). Managing climate change under uncertainty: Recursive integrated assessment at an inflection point. *Annual Review of Resource Economics 9*, 117–42.

Savage, L. J. (1954). *The foundations of statistics.* New York: Wiley.

Skiba, A. K. (1978, May). Optimal growth with a convex-concave production function. *Econometrica 46*(3), 527–539.

Stachurski, J. (2009). *Economic dynamics: theory and computation.* MIT Press.

Stern, N. (2007). *The Economics of Climate Change: The Stern Review.* Cambridge, UK: Cambridge University Press.

Sterner, T. and M. Persson (2008). An even sterner review: Introducing relative prices into the discounting debate. *Review of Environmental Economics and Policy 2*(1), 61–76.

Taconet, N., C. Guivarch, and A. Pottier (2021). Social cost of carbon under stochastic tipping points: when does risk play a role? *Environmental and Resource Economics 78*, 709–737.

The Climate Impact Lab (2023, September). *Documentation for Data-driven Spatial Climate Impact Model (DSCIM) Version092023-EPA.* The Climate Impact Lab.

Traeger, C. (2007). *Theoretical aspects of long-term evaluation in environmental economics*. Ph. D. thesis, University of Heidelberg. Published online at http://www.ub.uni-heidelberg.de/archiv/7049.

Traeger, C. (2010). Subjective risk, confidence, and ambiguity. *CUDARE working paper 1103*.

Traeger, C. (2014). Why uncertainty matters - discounting under intertemporal risk aversion and ambiguity. *Economic Theory 56*(3), 627–664.

Traeger, C. P. (2011). Discounting and confidence. *CUDARE working paper 1117*.

Traeger, C. P. (2012). A 4-stated dice: Quantitatively addressing uncertainty effects in climate change. *Environmental and Resource Economics 59*, 1–37.

Traeger, C. P. (2013). Discounting under uncertainty: disentangling the weitzman and the gollier effect. *Journal of Environmental Economics and Management 66*(3), 573–582.

Traeger, C. P. (2015). Closed-form integrated assessment and uncertainty. *CESifo Working Paper 5464*. Split into Traeger (2021) and Traeger (2023).

Traeger, C. P. (2019). Capturing intrinsic risk aversion. *SSRN Working Paper 3462905*.

Traeger, C. P. (2021). Uncertainty in the Analytic Climate Economy. *CEPR Discussion Paper DP16065*.

Traeger, C. P. (2022). IAMs and CO_2 emissions – An analytic discussion. Technical report. www.traeger.eu.

Traeger, C. P. (2023). ACE – Analytic Climate Economy. *American Economic Journal: Economic Policy 15*(3), 372–406.

Treasury, H. (2003). *The Green Book – Appraisal and Evaluation in Central Government*. London: HM Treasury.

Tsur, Y. and A. Zemel (1996). Accounting for global warming risks: Resource management under event uncertainty. *Journal of Economic Dynamics and Control 20*(6-7), 1289–1305.

Tsur, Y. and A. Zemel (2009). Endogenous discounting and climate policy. *Environmental and Resource Economics 44*, 507–520.

van den Bremer and van der Ploeg (2021). The risk-adjusted carbon price. *American Economic Review*, 2782–2810.

Van Den Bremer, T., C. Hambel, and R. v. der Ploeg (2023). Three reasons to price carbon under uncertainty: Accuracy of simple rules. *Available at SSRN 4666508*.

Van der Ploeg, F. and A. de Zeeuw (2018). Climate tipping and economic growth: precautionary capital and the price of carbon. *Journal of the European Economic Association 16*(5), 1577–1617.

Vissing-Jørgensen, A. and O. P. Attanasio (2003). Stock-market participation, intertemporal substitution, and risk-aversion. *The American Economic Review 93*(2), 383–391.

von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behaviour*. Princeton: Princeton University Press.

Wagener, F. O. O. (2003). Skiba points and heteroclinic bifurcations, with applications to the shallow lake system. *Journal of Economic Dynamics and Control 27*(9), 1533–1561.

Webster, M., A. P. Sokolov, J. M. Reilly, C. E. Forest, S. Paltsev, A. Schlosser, C. Wang, D. Kicklighter, M. Sarofim, J. Melillo, et al. (2012). Analysis of climate policy targets under uncertainty. *Climatic change 112*, 569–583.

Weil, P. (1990). Nonexpected utility in macroeconomics. *The Quarterly Journal of Economics 105*(1), 29–42.

Weitzman, M. (1974). Prices versus quantities. *Review of Economic Studies 41*(4), 477–91.

Weitzman, M. L. (1998). Why the far-distant future should be discounted at its lowest possible rate. *Journal of Environmental Economics and Management 36*, 201–208.

Weitzman, M. L. (2001). Gamma discounting. *The American Economic Review 91*(1), pp. 260–271.

Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics 91*(1), 1–19.